arXiv:1401.3020v4 [math.ST] 15 Oct 2015

EMPIRICAL GEODESIC GRAPHS AND CAT(K) METRICS FOR DATA ANALYSIS

KEI KOBAYASHI AND HENRY P WYNN

Abstract. A methodology is developed for data analysis based on empirically constructed geodesic metric spaces. For a probability distribution, the length along a path between two points can be defined as the amount of probability mass accumulated along the path. The geodesic, then, is the shortest such path and defines a geodesic metric. Such metrics are transformed in a number of ways to produce parametrised families of geodesic metric spaces, empirical versions of which allow computation of intrinsic means and associated measures of dispersion. These reveal properties of the data, based on geometry, such as those that are difficult to see from the raw Euclidean distances. Examples include clustering and classification. For certain parameter ranges, the spaces become CAT(0) spaces and the intrinsic means are unique. In one case, a minimal spanning tree of a graph based on the data becomes CAT(0). In another, a so-called "metric cone" construction allows extension to CAT(k) spaces. It is shown how to empirically tune the parameters of the metrics, making it possible to apply them to a number of real cases.

1. Introduction

There have been many development in statistics in which geometry, represented by an extension from Euclidean to more general spaces, has proved fundamental. Thus, reproducing kernel Hilbert spaces is a part of Gaussian process methods, Sobolev spaces are used for , and Besov spaces, for wavelets. Differential geometry, Riemannian manifolds, curvature and geodesics are at the core of information geometry, shape analysis and manifold learning.

This paper is concerned with CAT(0), and, more generally, CAT(k). Although related to Riemannian manifolds (k denotes the Riemannian curvature) these metric spaces have a less rigid structure and are qualitatively different from the spaces mentioned above. Trees are among the first examples in which CAT(0) properties have been used in statistics and bioinformatics [6]: the set of trees (the tree space) becomes a geodesic metric space (to be defined) when for each tree, the edge lengths are allocated to entries in a vector of sufficient dimension to capture all the tree structures. Then, the metric is the Euclidean geodesic distance on a simplicial complex called a "fan" which is induced by the tree structure: the geodesics are the shortest piecewise linear paths in the fan. Such spaces have already been shown to be CAT(0) [6] by Gromov's theory [15].

For a random variable X on a metric space $\mathcal{M}$ endowed with a metric $d(\cdot,\cdot)$ the general *intrinsic mean* is defined by

$$\mu = \arg\min_{m\in\mathcal{M}} \mathrm{E}\{d(X,m)^2\}.$$

Key words and phrases. intrinsic mean, extrinsic mean, CAT(0), curvature, metric cone, cluster analysis, non-parametric analysis.

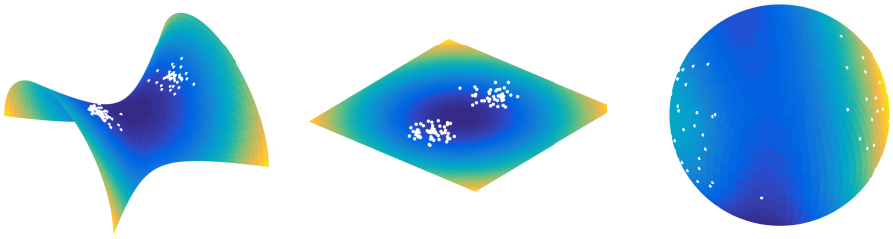

(a) Hyperboloid (b) Plane (c) Sphere

FIGURE 1. Function $f(m)$ for data on (a) a hyperboloid (curvature $c = -1$), (b) a plane ($c = 0$) and (c) a sphere ($c = 1$). The bluer represents the smaller value of $f(m)$. Only for the sphere, $f(m)$ has multiple minima. One hundred data points are sampled independently from a "Gaussian mixture," whose density is proportional to $\exp(-d(x, \mu_1)^2/2\sigma^2) + \exp(-d(x, \mu_2)^2/2\sigma^2)$ where $d(\mu_1, \mu_2) = \pi$ and $\sigma = \pi/8$.

The empirical intrinsic mean based on data $x = \{x_1, \ldots, x_n\}$, sometimes called the Fréchet mean, is defined as

$$\hat{\mu} = \arg\min_{m \in \mathcal{M}} f(m),$$

where

$$f(m) = \sum_{i=1}^{n} d(x_i, m)^2.$$

For Euclidean space, $\hat{\mu} = \bar{x}$, the sample mean. In general, $f(m)$ is not necessarily convex and the means, $\hat{\mu}$, are not unique. Figure 1 shows that the curvature can affect the property of $f(m)$. In particular, for so-called CAT(0) spaces, which (trivially) include Euclidean spaces, the intrinsic means $\hat{\mu}$ are unique.

Even when the mean is not unique, the function $f(m)$ can yield useful information, for example about clustering. We can also define second-order quantities:

$$s_0^2 = \inf_{m \in \mathcal{M}} \frac{1}{n} \sum_{i=1}^{n} d(x_i, m)^2 = \frac{1}{n} \sum_{i=1}^{n} d(x_i, \hat{\mu})^2$$

and

$$s_1^2 = \frac{2}{n(n-1)} \sum_{i<j}^{n} d(x_i, x_j)^2.$$

A key concept in the study of these issues is that the metrics are global geodesic metrics, that is metrics based on the shortest path between points measured by integration along a path with respect to a local metric. The interplay between the global and the local will concern us to a considerable extent.

For some intuition, consider a circle with arc length as distance and the uniform measure. The Fréchet intrinsic mean is the whole circle. We can say, informally, that the non-uniqueness arises because the space is not CAT(0). Although the curvature is zero in the Riemannian geometry sense, there are other considerations related to "diameter", which prevents it from having the CAT(0) property [17]. If we use the embedding Euclidean distance rather than the arc distance, we have

the *extrinsic mean*, but again, we obtain the whole perimeter. See [5] for further discussion on the intrinsic and extrinsic means. The empirical case is also problematic. For example, two data points on a diameter, with arc length as distance, give intrinsic means on the perpendicular diameter. Simple geometry shows that the empirical extrinsic mean is the whole circle. We also cover the more general CAT(k) spaces giving some new results related to "diameter" as well as conditions for the uniqueness of intrinsic means not requiring the spaces to be CAT(0).

The general form of $f(m)$ depends, here, on three parameters, α, β, γ, and it can be written in compact form:

$$f_{\alpha,\beta,\gamma}(m) = \sum_{i=1}^{n} \{g_\beta(d_\alpha(x_i, m))\}^\gamma,$$

where the function g_β and the construction of d_α are given below. Once we have introduced this new class of metrics, variety of statistics can be generalised; intrinsic mean, variance, clustering (based on local minima of $f(m)$). For classification problems, we can select an appropriate metric by cross-validation. Theoretically, we have the opportunity to apply well studies areas of geometry compared with methods based on selection of a good "loss function." In Table 1, we summarise such generalised statistics.

TABLE 1. A summary of generalised statistics by introducing α, β and γ.

	Euclidean	Generalised metric
Metrics	$d(x,y) = \|x-y\|$	$d_{\alpha\beta}(x,y) = g_\beta(d_\alpha(x,y))$
Intrinsic mean	$\arg\min_{m\in\mathbb{E}^d} \sum_{i=1}^n \|x_i - m\|^2$	$\arg\min_{m\in\mathcal{M}} \sum_{i=1}^n g_\beta(d_\alpha(x_i,m))^\gamma$
Variance	$\min_{m\in\mathbb{E}^d} \frac{1}{n}\sum_{i=1}^n \|x_i - m\|^2$	$\min_{m\in\mathcal{M}} \frac{1}{n}\sum_{i=1}^n g_\beta(d_\alpha(x_i,m))^\gamma$
Clustering function	$f(m) = \sum_{i=1}^n \|x_i - m\|^2$	$f_{\alpha\beta\gamma}(m) = \sum_{i=1}^n g_\beta(d_\alpha(x_i,m))^\gamma$

There are many ways to transform one metric into another, regardless of whether they are geodesic metrics. A straightforward way is to use a concave function g such that given a metric $d(\cdot,\cdot)$, the new metric is $d'(\cdot,\cdot) = g(d(\cdot,\cdot))$. This is plausible if we use non-convex $f_{\alpha,\beta,\gamma}$, which are useful, as will be explained, in clustering and classification. Such concave maps are often interpreted as loss functions, but we will consider them in terms of a change of metric which may lead to selection using geometric concepts. This is particularly true for the construction of g_β in this paper.

The basic definition and construction from a geodesic metric space to the special geodesics based on accumulation of density are given in the next section, together with the definition of a CAT(0) space. In Section 3, we first show that means and medians in simple one-dimensional statistics can be placed into our framework.

Because geodesics themselves are one-dimensional paths, this should provide some essential motivation. The d_α-metric is obtained by a local dilation. Our computational shortcut is to use *empirical* graphs, whose vertices are data points.

We will need, therefore, to define empirical geodesics. We start with a natural geodesic defined via a probability density function in which the distance along a path is the amount of density "accumulated" along that path. Then, an empirical version is defined whenever a density is estimated. We might have based the paper on kernel density estimates; instead, we have adopted a very geometric construction based on objects such as the the Delaunay complex and its skeleton graph.

In Section 4, the d_β metric is introduced. It is based on a function derived from a geodesic metric via shrinking, pointwise, to an abstract origin; that is to say an abstract cone is attached. The smaller the value of β, the closer to the origin (apex). The geometry of the two dimensional case drives our understanding of the general case because of the one-dimensional nature of geodesics. We prove that, for finite β, the embedding cone is CAT(k) with smaller k than the original space. We cover the more general CAT(k) spaces, giving some new results related to "diameter" , in Section 5, including conditions for the uniqueness of intrinsic means not requiring the spaces to be CAT(0).

Section 6 provides a summary of the effect of changing α and β. After some discussion on the selection of α and β in Section 7, Section 8 covers some examples.

2. Geodesics, intrinsic mean and extrinsic mean

The fundamental object in this paper is a geodesic metric space. This is defined in two stages. First, define a metric space $\mathcal{M} = (X, d)$ with base space X and metric $d(x, x')$. Sometimes, $\mathcal{M}$ will be a Euclidean space E_d of dimension d, containing the data points, but it may also be some special object such as a graph or manifold. Second, define the length of a (rectilinear) path between two points $x, x' \in X$ and the geodesic connecting x and x' as the shortest such path. This minimal length defines a metric $d^*(x, x')$, and the space endowed with the geodesic metric is called the geodesic metric space, $\mathcal{M}^* = \mathcal{M}(X, d^*)$.

The interplay between $\mathcal{M} = (X, d)$ and $\mathcal{M}^* = (X, d^*)$ will be critical for this paper, and, as mentioned, we will have a number of ways of constructing d^*.

For data points $x_1, \ldots, x_n$ in X, the empirical intrinsic (Fréchet) mean is

$$\mu = \arg \inf_{\mu \in X} \sum_{i=1}^{n} d^*(x_i, \mu)^2.$$

There are occasions when it is useful to represent $\mathcal{M}^*$ as a sub-manifold of a larger space (such as Euclidean space) $\mathcal{M}^+ = (X^+, d^+)$ with its own metric d^+. We can then talk about the extrinsic mean:

$$\mu^+ = \arg \inf_{\mu \in X} \sum_{i=1}^{n} d^+(x_i, \mu)^2.$$

Typically, the intrinsic mean is used as an alternative when the geodesic distance, d^* is hard to compute. The difficultly in considering the intrinsic mean in X^+ is that it may not lie in the original base space X. This leads to a third possibility, which is to project it back to X, in some way, as an approximation to the intrinsic mean μ (which may be hard to compute). We will discuss this again in Section 4.

2.1. **CAT(0) spaces.** CAT(0) spaces, which correspond to non-positive curvature Riemannian spaces, are important here because their intrinsic means are unique. The CAT(0) property is as follows. Take any three points $\{a, b, c\}$ in a geodesic metric space X and consider the "geodesic triangle" of the points based on the geodesic segments connecting them. Construct a triangle in Euclidean 2-space with vertices $\{a', b', c'\}$, called the comparison triangle, whose Euclidean distances, $\|a' - b'\|, \|b' - c'\|, \|a' - c'\|$, are the same as the corresponding geodesic distances just described: $d(a, b) = \|a' - b'\|$, etc. On the geodesic triangle select a point x on the geodesic edge between b and c and find the point x' on the edge $b'c'$ of the Euclidean triangle such that $d(b, x) = \|b' - x'\|$. Then the CAT(0) condition is that for all a, b, c and all choices of x:

$$d(x, a) \leq \|x' - a'\|.$$

For a CAT(0) space (i) there is a unique geodesic between any two points, (ii) the space is contractible, in the topological sense, to a point and (iii) the intrinsic mean in terms of the geodesic distance is unique. CAT(k) spaces, a generalization of CAT(0) spaces, are explained in Section 5.

2.2. **Geodesic metrics on distributions.** Let X be a d-dimensional Euclidean random variable absolutely continuous with respect to the Lebesgue measure, with density $f(x)$. Let $\Gamma = \{z(t), t \in [0, 1]\}$ be a parametrised integrable path between two points $x_0 = z(0), x_1 = z(1)$ in $\mathbb{R}^d$, which is rectifiable with respect to the Lebesgue measure. Let

$$s(t) = \sqrt{\sum_{i=1}^{d} \left(\frac{\partial z_i(t)}{\partial t} \right)^2},$$

with appropriate modification in the non-differentiable case, be the local element of length along Γ. The weighted distance along Γ is

$$d_\Gamma(x_0, x_1) = \int_0^1 s(t) f(z(t)) dt \tag{1}$$

The geodesic distance is

$$d(x_0, x_1) = \inf_\Gamma d_\Gamma(x_0, x_1).$$

Here we consider a random variable on Euclidean space but this can be generalized for Riemannian manifolds and even for singular spaces with a density with respect to a base measure naturally defined by the metric.

From the geodesic distances on distributions we shall follow three main directions:

(1) transform the geodesic metrics in various ways with parameters α, β to obtain a wide class of metrics,
(2) discover (locally) CAT(0) and CAT(k) spaces for certain ranges of the parameters,
(3) apply empirical versions of the metrics based on an empirical graph whose nodes are the data points.

There is an important distinction between global transformations applied to the whole distance between points and local transformations applied to dilate the distance element.

3. The d_α metric and minimal spanning trees

The general d_α metric is a dilation of the original distance d and what we have referred to as a local metric. It is obtained by transforming the density in (1). Thus for $\Gamma = \{z(t), t \in [0,1]\}$ between $x_0 = z(0)$ and $x_1 = z(1)$,

$$d_{\Gamma,\alpha}(x_0, x_1) = \int_0^1 s(t) f^\alpha(z(t)) dt$$

and

$$d_\alpha(x_0, x_1) = \inf_\Gamma d_{\Gamma,\alpha}(x_0, x_1).$$

Changing α essentially changes the local curvature. Roughly speaking, when α is more negative (positive), the curvature is more negative (positive).

In the next subsection, we look at the one-dimensional case. Although this case is elementary, good intuition is obtained by rewriting the standard version in terms of a geodesic metric.

3.1. **One-dimensional means and medians.** Assume that X is a continuous univariate random variable with probability density function $f(x)$ and cumulative distribution function (CDF) $F(x)$. The mean $\mu = \mathrm{E}(X)$ achieves $\min_m \mathrm{E}\{(X-m)^2\}$. Here we are using the Euclidean distance: $d_E(x,y) = |x-y|$.

The median is defined by $\nu = F^{-1}(1/2)$. On a geometric basis, we can say that ν achieves $\min_m d_D(m,x)^2$, where we use a metric that measures the amount of probability between x and z:

$$d_D(x,z) = |F(x) - F(z)|. \tag{2}$$

Carrying out the calculations:

$$\begin{aligned} \mathrm{E}_X\{d_D(m,X)^2\} &= \int_{-\infty}^{\infty} (F(m) - F(x))^2 f_X(x) dx \\ &= \frac{1}{3} - F(m)(1 - F(m)) \end{aligned}$$

which achieves a minimum of $\frac{1}{12}$ at $F(m) = \frac{1}{2}$, as expected.

Now let us consider the sample version. Let $x_{(1)}, \ldots, x_{(n)}$ be the order statistics, which we assume are distinct. One of the first exercises in statistics is to show that $\hat{\mu} = \frac{1}{n}\sum_i x_i = \bar{x}$ minimises $\sum (x_i - m)^2$, with respect to m.

For the median, first consider using the first of the two approaches with the empirical CDF $\hat{F}(x)$. We obtain various definitions depending on our definition of F and $\hat{F}^{-1}$, or just using convention. Using the metric approach the natural metric is to take

$$\tilde{d}_1(x,z) = \mid \hat{F}_D(x) - \hat{F}(z) \mid,$$

with the standard definition of $\hat{f}$. Applied to distinct data points x_i, x_j this is equal to $\mid i - j \mid$. For an arbitrary m

$$\begin{aligned} \tilde{d}_1(x_i, m) &= \mid \hat{F}(x_i) - \hat{F}(m) \mid \\ &= \frac{1}{n} \mid i - i(m) \mid, \end{aligned}$$

where $i(m) = \max\{i : x_{(i)} \leq m\}$. Then, $\min_m \sum_{i=1}^{n} (i - i(m))^2$ is achieved at $x_{(\frac{n+1}{2})}$ when n is odd and at $x_{(\frac{n+2}{2})}$ when n is even.

Another approach for the median would be to take a piecewise linear approximation to F which is equivalent to having a density $\hat{f}$ that is proportional to $\frac{1}{x_{(n+1)} - x_{(n)}}$ in the interval $[x_{(n)}, x_{(n+1)})$. Then, the metric is

$$\tilde{d}_2(x, z) = \int_{\min(x,z)}^{\max(x,z)} \hat{f}(y) dy,$$

and $\min_m \sum_{i=1}^{n} \tilde{d}_2(x_i, m)^2$ is achieved at $x_{(\frac{n+1}{2})}$ when n is odd and at $\frac{1}{2}(x_{(\frac{n}{2})} + x_{(\frac{n+2}{2})})$, when n is even. We can think of this last result in another way. Consider the points $y_i = \hat{F}(x_{(i)}) = \frac{i}{n}$ as points in $[0, 1]$, take the empirical mean of the points and transform them back with the CDF corresponding to $\hat{f}$, namely, the piecewise linear approximation.

The idea of weighting intervals should provide intuition when we extend the intervals to edges on a graph, because edges are one-dimensional.

3.2. **The d_α metric for graphs.** There are a number of options to produce an empirical version of the d_α metric, based on the data. One such option would be to produce a smooth empirical density $f(t)$ followed by numerical integration and optimization to compute the geodesics. We prefer a much simpler method based on a graph constructed from the data. All geodesic computation is then restricted to the graph. We list some candidates (see appendix A.1 for a description of the second and third candidates listed below):

(1) the complete graph with vertices at the data points and all edges,
(2) the edge graph (1-skeleton) of the Delaunay simplicial complex with vertices at the data points,
(3) the Gabriel graph with vertices at the data points.

The discussion below applies to the complete graph or any connected sub-graph. For any such graph, define a version of the d_α distance just for edges,

$$\tilde{d}_{\alpha,ij} = d_{ij}^{1-\alpha},$$

where d_{ij} is the Euclidean distance from x_i to x_j. This can be explained by making a transformation

$$ds \to \frac{ds}{d_{ij}}.$$

We refer to this as *edge regularization*. We then apply α in the usual way to obtain

$$\frac{ds}{d_{ij}^{\alpha}}.$$

The new "length" of each edge e_{ij} is obtained by integrating this "density" along the edge. In this sense, d_{ij} also plays the role of density estimation. Although we need a regularization $d_{ij}^{-1/p}$ with respect to the dimension p for density estimation [22], we manage the regularization by rescaling the parameter α. Note that $\alpha = 1$ gives the unit length and $\alpha = 0$ restores the original length.

Now we consider only the set of edges E of the graph $G(V, E)$ as a metric space with the metric defined by the geodesic:

$$\tilde{d}_\alpha(x_0, x_1) = \inf_\Gamma \sum_{(i,j)\in\Gamma} \tilde{d}_{\alpha,ij},$$

where the infimum is taken over all (connected) paths Γ between x_0 and x_1. Note that even if $d_{ij}^{-1/p}$ can estimate the local density well, it does not follow that the metric d_α can be approximate by the metric $\tilde{d}_\alpha$ since edge lengths d_{ij} on each path Γ are not independent. It is suggested that further theoretical work is necessary. Here we will admit $\tilde{d}_\alpha$ as an approximation of d_α.

If the graph is not a complete Euclidean graph with weights equal to the Euclidean lengths of the edges, some edges may not be in any edge geodesics between any pair of vertices.

Definition 3.1. *For an edge-weighted graph G with weights $\{d_{ij}\}$ on the graph, G^*, which is the union of all the edge geodesics between all pairs of vertices, is called the geodesic sub-graph (or geodesic graph) of G.*

We will see how the geodesic sub-graphs transform as the value of α changes.

We make an important general position assumption that the set of values $\{d_{ij} \mid (i,j) \in E\}$ are distinct, that is there are no ties. This is an additional general position assumption to that given for the Delaunay cells (see appendix A.1). We order the values using only a single suffix for simplicity: $d_1 < d_2 < \cdots < d_M$ where $M = |E|$. For $\alpha < 1$, this induces the $\tilde{d}_{\alpha,i} (= d_i^{1-\alpha})$ values:

$$\tilde{d}_{\alpha,1} < \tilde{d}_{\alpha,2} < \cdots < \tilde{d}_{\alpha,M}.$$

Now, consider the geodesics as $\alpha \to -\infty$. Recall that a circuit in a graph is a connected path that begins and ends in some vertex and an elementary circuit is a circuit that visits a vertex no more than once. Consider an edge $(i,j) \in E$ that has the following property which we call Q: it is in an elementary circuit $\mathcal{C}$ of the graph in which all other edges have smaller values of d_{ij} namely

$$d_{rs} < d_{ij} \text{ for } (r,s) \in \mathcal{C},\ (r,s) \neq (i,j).$$

Then, the path $\Gamma(i,j)$ (within the circuit) from x_i to x_j not containing the edge (i,j) has length smaller than $\tilde{d}_{\alpha,ij}$ when α is sufficiently negative:

$$\sum_{(r,s)\in\Gamma(i,j)} d_{rs}^{1-\alpha} < d_{ij}^{1-\alpha}$$

From this argument, we see that for sufficiently large $|\alpha|$ as α approaches $-\infty$, every edge having property Q is removed from the geodesic sub-graph, and we obtain a tree.

Let us summarize this algorithm, which applies to a general edge-weighted graph with distinct edges. We refer to this algorithm as the *backwards* algorithm. It clearly gives a tree.

(1) Let $|E| = M$ and label the edges $e_1, \ldots, e_M$ in increasing order of their weights.
(2) Starting with edge e_M, remove e_M if it is in a cycle otherwise continue to e_{M-1}.
(3) (General step) Continue downwards at each stage removing an edge if it is in a cycle of the remaining subgraph.

(4) Stop if no more edges can be removed using step 3.

There is a natural *forwards* algorithm that also yields a tree as follows.

(1) Let $|E| = M$ and label the edges $e_1, \ldots, e_M$ in increasing order of their weights.
(2) Starting with e_1, add an edge if adding it does not create a cycle.
(3) (General step) Continue adding an edge at each step provided that the addition does not create a cycle.
(4) Stop if no more edges can be added.

We have the following.

Lemma 3.2. *Given a connected edge-weighted graph $G(V, E)$, the backward and forward algorithms yield the same tree, which we call $T^*(G)$.*

Proof. Let T_1 and T_2 be the trees generated by the backward and forward algorithms respectively. Any edge of G not in T_1 cannot be in T_2, because it is in a circuit of edges with lower weights than itself, by the forward construction. This is sufficient because each tree has the same number of edges. In fact, the tree can be constructed by the simple rule: remove all edges that are in a circuit with "smaller" edges; the order in which the edges are removed is irrelevant. □

Although there are several implementations of forward and backward minimal spanning tree algorithms, their importance here is to give intuition for the construction of geodesics as $\alpha \to -\infty$.

We now show that the tree $T^*(G)$ is the minimal spanning tree in a strong sense.

Theorem 3.3. *Let $G(V, E)$ be a graph with $|V| = n$ and distinct edge weights $\{d_{ij},\ (i,j) \in E\}$. There is a unique spanning tree $T^*(G)$ whose ordered weights $d_1^* < d_2^* < \cdots < d_{n-1}^*$ have the following minimal property. If $d_1' < d_2' < \cdots < d_{n-1}'$ are the ordered weights of any other spanning tree, then*

$$d_i^* \le d_i',\ i = 1, \ldots, n-1,$$

with strict inequality for at least one $i = 1, \ldots, n-1$. Moreover $T^*(G)$ is given by the forward (or backward) algorithm.

Proof. The proof is obtained by contradiction. Let T^* be the tree constructed by the forward algorithm. Let T' be another spanning tree with ordered weights $d_1' < d_2' < \cdots < d_{n-1}'$ and suppose that for some j with $1 \le j \le n-1$, we have $d_j^* > d_j'$. Let the edges of T^* and T' be $e_{(1)}^* < \cdots < e_{(n-1)}^*$ and $e_{(1)}' < \cdots < e_{(n-1)}'$, respectively. Since $d_{(j)}^* > d_{(j)}'$, we must have $e_{(j)}^* > e_{(j)}'$. In the sequence $e_1, \ldots, e_M$, let $e_{(j)}^* = e_r$ and $e_{(j)}' = e_s$; then, we must have $s < r$. By the nature of the forward algorithm, new vertices are used only when an edge is included by the algorithm: all unused edges form circuits and therefore use edges of the tree constructed up to that point. Thus, the subgraph of $G(V, E)$ with edges $e_1, \ldots, e_r$ has exactly $j+1$ vertices; this value is attained on the addition of the last edge e_r. However, T', which is a subgraph of $G(V, E)$, attains $j+1$ vertices at e_s with $s < r$ which is a contradiction. □

The ordering property in Theorem 3.3 is known from the theory of stochastic ordering: if the empirical CDFs of the distances $\{d_i^*\}$ and $\{d_i'\}$ are G^* and G', respectively, then $G^* \ge G'$ with strict inequality for at least one value (in fact, over at least one non-empty interval). Thus, not only $\sum_i d_i^* < \sum_i d_i'$ but also $\sum_i g(d_i^*) < \sum_i g(d_i')$ for any non-decreasing function g.

For sufficiently negative α, the tree T itself, that is the tree as a metric space with metric d_α, is a CAT(0) space. We need to extend the metric somewhat so that it applies to the edges, in addition to the nodes. Thus, for any two points x, x' on the tree, define

$$d_\alpha(x, x') = \inf_{\Gamma(x,x')} \int_{\Gamma(x,x')} w(s)ds,$$

where the integral is taken along the (unique) path $\Gamma(x, x')$ on the tree and $w(s) = \frac{1}{d_{ij}^\alpha}$ when line element ds is in edge e in $\Gamma(x, x')$.

Theorem 3.4. *There is an α^* such that for any $\alpha \leq \alpha^*$, the geodesic sub-graph becomes the minimal spanning tree $T^*(G)$ endowed with the d_α metric and, therefore, becomes a CAT(0) space.*

We see that for sufficiently negative α, *every* geodesic defined with the d_α metric lies in the tree T^*. In fact, although we started with a general connected graph, any graph for which the edges can be mapped into a Euclidean interval gives a CAT(0) tree using this construction.

There are some well-known algorithms and considerable literature on minimal spanning trees. Remarkably, the minimal spanning tree for a complete Euclidean graph is the same as the minimal spanning tree for the Delaunay graph and is therefore a subgraph of the Delaunay graph.

Theorem 3.5. *Let G be a complete edge-weighted graph whose weights are distinct and monotonically related to those of the complete Euclidean graph G_E of a point set X in d dimensions. Then, G, G_E, the Gabriel graph and the Delaunay graph of the point set X all have the same and unique minimal spanning tree.*

Versions of this result are known in two dimensions [2], but the authors had some difficulty in finding a concise proof in the literature. We present a proof in appendix A.3.

Thus, any algorithm for finding the minimal spanning tree of a graph will give the minimal spanning tree of the Delaunay graph, without having to find the full Delaunay complex for any Euclidean space into which the points can be embedded. Further, as seen previously, the minimal spanning tree has the strong minimal property in the sense of Theorem 3.5.

3.3. **The double α-chain.** We may study the geometry as α increases away from $-\infty$. Following Theorem 3.3, we are interested in two cases: the Euclidean Delaunay graph and the complete Euclidean graph. In both cases, we consider the d_α metric.

Theorem 3.6. *Let G_α be an edge-weighted graph with distinct weights $\{d_{ij}^{1-\alpha}\}$ and let G_α^* be its geodesic subgraph; then,*

$$|1 - \alpha'| > |1 - \alpha| \Rightarrow G_{\alpha'}^* \subseteq G_\alpha^*$$

Proof. This follows from the consideration of geodesics. An edge (i, j) in G is not in G_α^* if it is not a geodesic. In this case, there is an alternative path Γ from i to j such that $d_{ij}^{1-\alpha} > \sum_{(r,s)\in\Gamma} d_{rs}^{1-\alpha}$. However, this inequality is preserved if α is decreased, so that $1 - \alpha$ is increased. Thus an edge absent from G_α^* is absent from $G_{\alpha'}^*$. □

3.4. α **and CAT(k).** Let $C(X,p,r) := \{x \in X \mid d(p,x) \leq r\}$ be a geodesic disk of radius $r \geq 0$ centred at $p \in X$. Define the maximum radius $D_k(X,x)$ of the disk centred at x as being CAT(k), that is

$$D_k(X,x) := \sup\{r \geq 0 \mid X \cap C(X,x,r) \text{ is CAT(k)}\}.$$

See Section 5 for the definition of CAT(k). If X is a metric graph, $D_k(X,x) = \sup\{r \geq 0 \mid X \cap C(X,x,r)$ does not include a cycle whose length is less than $2\pi/\sqrt{\max(k,0)}\}$.

Consider a rescaling of X such that the shortest (longest) edge length is 1, and denote it as $\bar{X}$ for $\alpha \leq 1$ ($\alpha > 1$).

Theorem 3.7. *If* $|\alpha' - 1| > |\alpha - 1|$,

$$D_k(\bar{G}^*_{\alpha'}, x) > D_k(\bar{G}^*_{\alpha}, x) \text{ for each } k \in \mathbb{R}.$$

Proof. Because the α-chain is increasing for $\alpha < 1$, each cycle in $\bar{G}^*_\alpha$ is removed one by one as α decreases. Furthermore, each cycle length increases as α decreases because, by the rescaling, every edge length is greater than 1 and it increases as α decreases. This gives the decreasing property of $D_k(\bar{G}^*_\alpha, x)$ for $\alpha \leq 1$. We can prove the result for $\alpha > 1$ similarly. □

By the theorem, $\bar{G}^*_\alpha$ becomes "more CAT(k)" for a smaller $\alpha < 1$. Because rescaling of the graph does not affect the uniqueness of the intrinsic mean, G^*_α tends to have a unique mean for a smaller $\alpha < 1$.

3.5. **Geodesic graphs in 2-d with different** α**.** Figures 2 (a)-(f) are geodesic graphs with different values of α for 50 samples of the standard 2-d Normal distribution. We give two cases in which we decrease α: the Delaunay graph in Figure 2 and the complete graph in Figure 3. By the time $\alpha = -0.3$ the cases are indistinguishable and have the same minimal spanning geodesic graph for large negative values of α, as expected.

This is predictable from Theorem 3.5 and 3.6 and gives an important practical strategy: when the dimension is high and α is small, use the complete graph rather than the Delaunay graph because the former requires computational cost only proportional to d, whereas the computational cost of the latter is $O(n^{d/2})$ [10] [8].

4. The d_β metric

The CAT(0) property of Euclidean space implies that we do not obtain multiple local minima of f even for multi-modal distributions. However, an appropriate concave transformation of the metric can modify the base data space making it less CAT(0). We introduce the d_β metric via a transformation g_β as a candidate.

For any geodesic metric space (X,d) with metric $d(x_0,x_1)$ and a parameter $\beta > 0$, we can define the metric

$$d_\beta(x_0,x_1) = g_\beta(d(x_0,x_1))$$

where

$$g_\beta(z) = \begin{cases} \sin(\frac{\pi z}{2\beta}), \text{ for } 0 \leq z \leq \beta, \\ \\ 1, \text{ for } z > \beta. \end{cases}$$

Since g_β is a concave function on $[0,\infty)$ (see Figure 4), d_β becomes a metric but

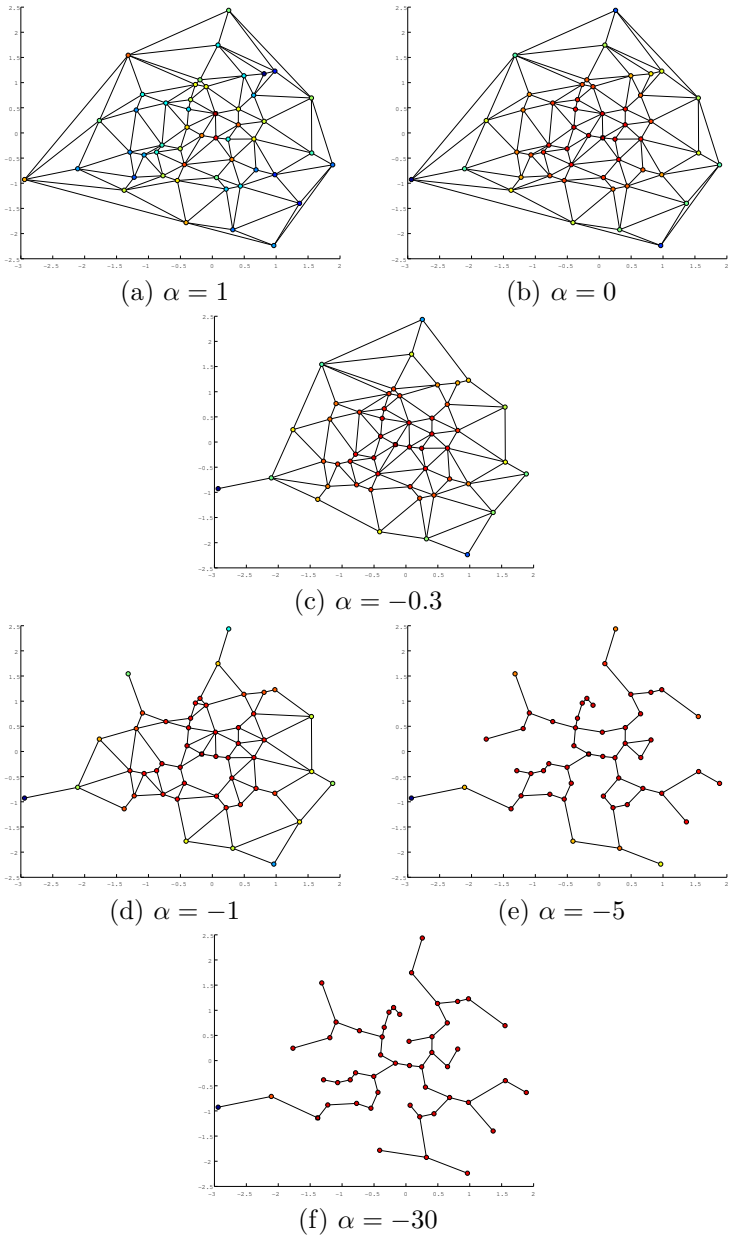

(a) $\alpha = 1$ (b) $\alpha = 0$

(c) $\alpha = -0.3$

(d) $\alpha = -1$ (e) $\alpha = -5$

(f) $\alpha = -30$

FIGURE 2. Geodesic graphs with different values of α for 50 samples from the standard 2-d Normal distribution. The initial graph ($\alpha = 0$) is the Delaunay graph. The value $f(x) = \sum_i d_\alpha(x_i, x)^2$ for each sample point x is represented by the colours red (small) and blue (large), and the minimum is represented by a square.

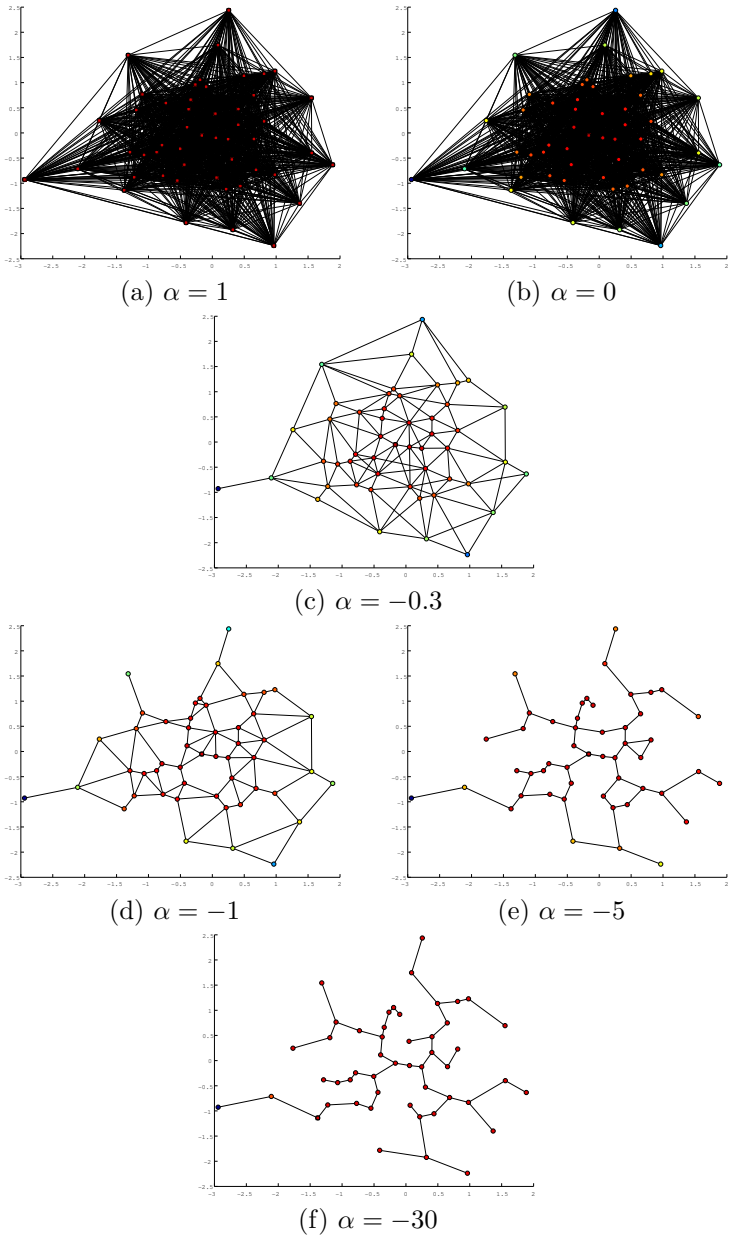

(a) $\alpha = 1$ (b) $\alpha = 0$

(c) $\alpha = -0.3$

(d) $\alpha = -1$ (e) $\alpha = -5$

(f) $\alpha = -30$

FIGURE 3. Geodesic graphs with different values of α for 50 samples from the standard 2-d Normal distribution. The initial graph ($\alpha = 0$) is the complete graph. The value $f(x) = \sum_i d_\alpha(x_i, x)^2$ for each sample point x is represented by the colours red (small) and blue (large), and the minimum is represented by a square.

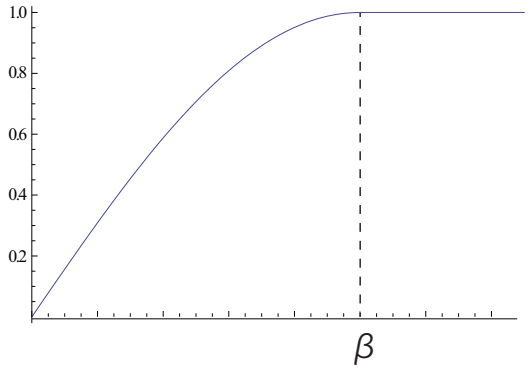

FIGURE 4. Graph of g_β for a finite positive β.

not necessarily a geodesic metric. We can express this conveniently as $g_\beta(z) = \max\left(\sin(\frac{\pi z}{2\beta}), H(z-\beta)\right)$, where H is the Heaviside function,

$$H(z) = \begin{cases} 1, & \text{for } x \geq 0, \\ \\ 0 & \text{otherwise.} \end{cases}$$

It is easiest to consider the case that $d(x_0, x_1)$ is the Euclidean distance on the real line. As $\beta \to \infty$ for small values of d, the metric behaves like $\frac{\pi d}{2\beta}$, and as $\beta \to 0$, d behaves like $\sin(d)$ rescaled to $(0, \beta]$. For Euclidean distances greater than β, d_β returns a constant distance of unity. The metric has the effect of downsizing large distances to unity. Because, as will soon be seen, d_β can be recognized as a geodesic metric of a cone embedding X, we refer to the mean

$$\hat{\mu}_\beta = \arg \inf_{\mu \in X} \sum_{i=1}^{n} g_\beta(d(x_i, \mu))^2$$

as the *β-extrinsic mean*.

We consider a scheme in which the real line, or part of it, is mapped into the unit circle, where it can be represented by an angle θ. In the unit disk, D, a point is represented by polar coordinates: $(y_1, y_2) = (r\cos(\theta), r\sin(\theta))$.

We consider the top half of the unit disk, for which $y_2 \geq 0$, namely

$$D^+ = \{(r, \theta) : 0 \leq r \leq 1,\ 0 \leq \theta \leq \pi\}.$$

We will now give a rule for travelling from a point $y = (y_1, y_2)$ in D to the point P:$(1, 0)$.

(C1) If $y \in D^+$, travel in a straight line to P. This is the geodesic for the Euclidean metric

(C2) If $\theta > \pi$, travel first in straight line to the origin $(0, 0)$ and then in a straight line from $(0, 0)$ to P.

We now apply a similar rule to points $y = (r\cos(\theta), r\sin(\theta))$ and $y' = (r'\cos(\theta'), r\sin(\theta'))$.

(D1) If $\mid \theta - \theta' \mid \leq \pi$, we take the Euclidean distance

$$\sqrt{r^2 + r'^2 - 2rr'\cos(\mid \theta - \theta' \mid)}$$

(D2) If $\mid \theta - \theta' \mid > \pi$, we take $r + r'$, indicating a route to the origin and out to the other point.

Now, consider two points at θ and θ' on the circumference of the circle; then, in case (D1) above, the distance between them is

$$\sqrt{2 - 2\cos(\mid \theta - \theta' \mid)} = 2\sin\left(\frac{\mid \theta - \theta' \mid}{2}\right).$$

When $\mid \theta - \theta' \mid > \pi$, we obtain the value 2.

Now, keeping D^+ fixed, let us rescale the distance inside the argument. In terms of the two points just described, we place the first point at P so that $\theta = 0$ and scale the second with a factor $\frac{1}{\beta}$. Then, without changing the notation, and removing the factor of 2, we have $\sin(\frac{\pi\theta}{2\beta})$. Then, $\beta > 1$ implies that we may use the Euclidean distance for a wider interval. Case (D2) above corresponds to the distance 2 before rescaling and the new version $\sin(\frac{\pi\theta}{2})$ achieves the value 1. This corresponds precisely to the d_β metric of this section.

4.1. **The β-extrinsic mean: one dimension.** The extrinsic mean keeps the mean within the original data space but uses the d_β metric. We will see below that the d_β metric involves, even in the general case, an extension of the original space by a single dimension. This distinguishes it from the embeddings for extrinsic means used in the literature, which involve high-dimensional Euclidean space, for example the case of the tree space.

Controlling β, as will be seen below, controls the value of k when the embedding space is considered as a CAT(k) space. We have an indirect link between clustering and CAT(k) spaces. As β decreases while the embedding space becomes more CAT(0) (k decreasing) the original space becomes less CAT(0). This demonstrates, we believe, the importance of the CAT(k) property in geodesic-based clustering.

In Euclidean space, the standard Euclidean distance dose not exhibit multiple local means because the space is trivially CAT(0). However, by using the d_β-metric with a sufficiently small β, the space can have multiple local means, as shown in Figure 5.

Figure 6 shows a plot of the local minima of $f(x) = \sum_i d_\beta(x_i, x)^2$ against β for the same samples as those in Figure 5.

4.2. **The general case: metric cone.** The above construction is a special case of a general construction that applies to any geodesic metric space and hence to those in this paper. Let $\mathcal{X}$ be a geodesic metric space with a metric $d_\mathcal{X}$. A metric cone $\tilde{\mathcal{X}}_\beta$ with $\beta \in (0, \infty)$ is a cone $\mathcal{X} \times [0, 1]/\mathcal{X} \times \{0\}$ with a metric

$$\tilde{d}_\beta((x, s), (y, t)) = \sqrt{t^2 + s^2 - 2ts\cos(\pi \min(d_\mathcal{X}(x, y)/\beta, 1))}$$

for any $(x, s), (y, t) \in \tilde{\mathcal{X}}_\beta$.

The intuitive explanation is as follows. See Figure 7. Let $\mathcal{X}_\beta$ be a subset $\{(x, 1) \mid (x, t) \in \tilde{\mathcal{X}}_\beta\}$ with the geodesic metric on $\mathcal{X}_\beta$. Thus, $\mathcal{X}_\beta$ and $\mathcal{X}$ are the same as a set but endowing different metrics. Then $\mathcal{X}_\beta$ is a rescaling of the metric on $\mathcal{X}$ by β. For any $(x, s), (y, t) \in \tilde{\mathcal{X}}_\beta$, their projections $(x, 1), (y, 1)$ give two points $x, y \in \mathcal{X}$, respectively. For a geodesic $\gamma \subset \mathcal{X}$ between x and y, consider a cone $\{(z, s) \mid z \in \gamma, s \in [0, 1]\}$ spanned by γ. This cone can be isometrically embedded into an "extended unit circular sector", i.e. a covering $\{(r, \theta) \mid r \in [0, 1], \theta \in (-\infty, \infty)\}/\{(0, \theta) \mid \theta \in (-\infty, \infty)\}$ of the unit disk corresponding to

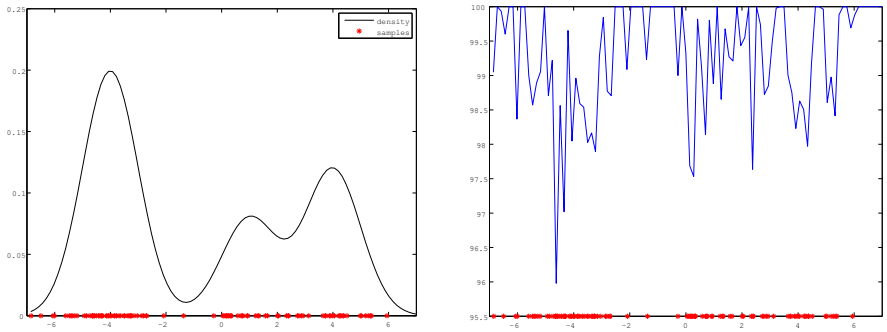

(a) Density function (b) $\beta = 0.1$

(c) $\beta = 1$ (d) $\beta = 10$

FIGURE 5. (a) The density function is a mixture of three normal distributions and 100 i.i.d. samples from it. (b)-(d) Graphs of $\sum_i d_\beta(x_i, m)^2$ against m for different values of β.

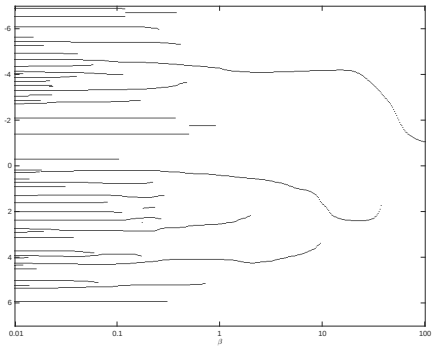

FIGURE 6. Plot of the local minima of $\sum_i d_\beta(x_i, m)^2$ against β.

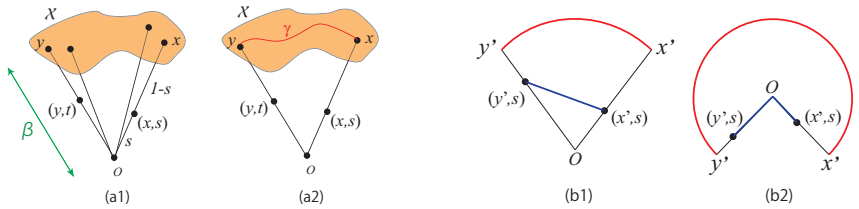

FIGURE 7. How to define a metric cone for a geodesic metric space $\mathcal{X}$.

$\theta \in [0, \pi d_{\mathcal{X}}(x,y)/\beta]$. Then (x,s) and (y,t) are also mapped into the extended unit circular sector; the distance $\tilde{d}_\beta((x,s),(y,t))$ for $\beta = 1$ corresponds to the case (D2) of a disk if we set $(r,r') = (s,t)$ and $(\theta,\theta') = (\pi x, \pi y)$. This corresponds to the length of the blue line path in Figure 7 (b1) and (b2). For further details on metric cones, refer to [11].

The following result indicates that the metric cone space preserves the CAT(0) property of the original space and the smaller values of β continue this process. See section 6.3 for more details.

Theorem 4.1. (1) *If $\mathcal{X}$ is a CAT(0) space, the metric cone $\tilde{\mathcal{X}}_\beta$ is also CAT(0) for every $\beta \in (0,\infty)$.*
(2) *If $\tilde{\mathcal{X}}_{\beta_2}$ is CAT(0), $\tilde{\mathcal{X}}_{\beta_1}$ is also CAT(0) for $\beta_1 < \beta_2$.*
(3) *If $\mathcal{X}$ is CAT(k) for $k \geq 0$, $\tilde{\mathcal{X}}_\beta$ becomes CAT(0) for $\beta \leq \pi/\sqrt{k}$.*

The proof is given in appendix A.2.

5. CAT(k) SPACES, CURVATURE, DIAMETER AND EXISTENCE OF MEANS

A CAT(k) space for $k \in \mathbb{R}$ is a geodesic metric space satisfying the following CAT(k) condition. Take any geodesic triangle Δabc whose perimeter is less than $2\pi/\sqrt{\max(k,0)}$ and select any point x on a geodesic $\widetilde{bc}$. Let $\Delta a'b'c'$ be a comparison triangle, which has the same edge length as Δabc, on a surface with a constant curvature k, i.e. a sphere with radius $1/\sqrt{k}$ for $k > 0$, a plane for $k = 0$ and a hyperbolic space for $k < 0$. Set x' on $\widetilde{b'c'}$ satisfying $|\widetilde{bx}| = |\widetilde{b'x'}|$. Then the CAT($k$) condition is that for all a, b, c and all choices of x, $|\widetilde{ax}| \leq |\widetilde{a'x'}|$. Thus every CAT($k$) space is a CAT($k'$) space for $k < k'$. Every metric graph is CAT(k) for $k > 0$ if and only if there is no loop shorter than $2\pi/\sqrt{k}$ because every metric tree is CAT(0) and therefore CAT(k) for $k > 0$.

Let $\mathcal{X}$ be a geodesic metric space and fix it throughout this section. The diameter of a subset $A \subset \mathcal{X}$ is defined as the length of the longest geodesic in A. We define classes $\mathcal{C}_{\text{convex}}$, $\mathcal{C}_{L_\gamma}$ and $\mathcal{C}_{\text{geodesic}}$ as follows.

(1) $\mathcal{C}_{\text{convex}}$: the class of subsets $A \subset \mathcal{X}$ such that the geodesic distance function $f_p(x) := d(p,x)$ is strictly convex on A for each $p \in A$. Here, "convex" means geodesic convex, i.e. a function f on $\mathcal{X}$ is convex iff for every geodesic $\{\gamma(t) \mid t \in (t_0,t_1)\}$ on $\mathcal{X}$, $f(\gamma(t))$ is convex with respect to t.
(2) $\mathcal{C}_{L_\gamma}$ for $\gamma \in [1,\infty]$: the class of the subsets $A \subset \mathcal{X}$ such that for any probability measure whose support is in A and non-empty, the intrinsic L_γ-mean exists uniquely. We refer to $\mathcal{C}_{L_2}$ as $\mathcal{C}_{\text{mean}}$.

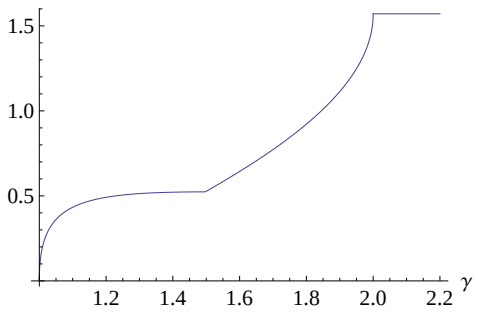

FIGURE 8. Graph of $\theta_0(\gamma)$.

(3) $\mathcal{C}_{\text{geodesic}}$: the class of subsets $A \subset \mathcal{X}$ such that for every pair $p, q \in A$, the geodesic between p and q is unique.

Lemma 5.1.

$$\mathcal{C}_{\text{convex}} \subset \mathcal{C}_{L_\gamma} \subset \mathcal{C}_{\text{geodesic}}$$

for any $\gamma \in [1, \infty]$.

Proof. If $A \in \mathcal{C}_{\text{convex}}$, $f_y(x) = d(y, x)$ is a strictly convex function on A for each $y \in A$; hence, $\int d(y, x)\mathrm{d}\mu_Y$ is strictly convex for any probability measure μ whose support is in A and non-empty. Thus, $A \in \mathcal{C}_{L_\gamma}$. Next, assume that $B \notin \mathcal{C}_{\text{geodesic}}$ and $x, y \in B$; then, there are at least two different geodesics, γ_1 and γ_2, between x and y. Thus, there are two points x' and y' in $\gamma_1 \cap \gamma_2$ such that there is no intersection of γ_1 and γ_2 between x' and y'. Then, the mid points of x' and y' on each geodesic become intrinsic L_γ-means of the measure with two equal point masses on x' and y'. This implies that $B \notin \mathcal{C}_{L_\gamma}$. □

Let $D_{\text{convex}}, D_{L_\gamma}$ and D_{geodesic} be the largest values (including ∞) such that every subset whose diameter is less than the value belongs to $\mathcal{C}_{\text{convex}}, \mathcal{C}_{L_\gamma}$ and $\mathcal{C}_{\text{geodesic}}$, respectively. Then, evidently from Lemma 5.1, $D_{\text{convex}} \leq D_{L_\gamma} \leq D_{\text{geodesic}}$ for $1 \leq \gamma \leq \infty$.

Note that if $\mathcal{X}$ is CAT(0), $D_{\text{convex}} = D_{L_\gamma} = D_{\text{geodesic}} = \infty$. In general, the following theorem holds.

Theorem 5.2. (1) *If* $\mathcal{X}$ *is CAT(k),* $D_{\text{convex}} \geq \pi/(2\sqrt{\max(k, 0)})$.
(2) *If* $\mathcal{X}$ *is CAT(k),* $D_{\text{geodesic}} \geq \pi/\sqrt{\max(k, 0)}$.
(3) *If* $\mathcal{X}$ *is a surface with a constant curvature* $k > 0$, $D_{L_1} \geq \pi/(2\sqrt{k})$.

The proof is given in appendix A.4. By Theorem 5.2(1), $D_{L_\gamma} \geq D_{\text{convex}} \geq \pi/(2\sqrt{k})$. Thus, a lower curvature k gives a wider area where the intrinsic L_γ-mean is unique. According to Theorem 5.2(3), this lower bound for D_{L_1} is the best universal upper bound for any $\mathcal{X}$ with CAT(k) property.

For $\gamma > 1$, D_{L_γ} is bounded above by $(\theta_0(\gamma) + \pi/2)/\sqrt{k}$ where θ_0 is an increasing function of γ, as proved in appendix A.5. This upper bound shows that the parameter γ plays a role in controlling the uniqueness of the mean, but it does not in Euclidean space, where the L_γ-mean functions are always convex.

6. The α, β, γ: a summary

Combining the three deformations by α, β and γ, a novel Fréchet mean is proposed:

$$\hat{\mu} = \arg\min_{m \in \mathcal{M}} f(m),$$

where

$$f(m) = f_{\alpha,\beta,\gamma}(m) = \sum_i \{g_\beta(d_\alpha(x_i, m))\}^\gamma.$$

We now review three main areas: (i) CAT(0) and CAT(k) aspects, (ii) uniqueness of means and (iii) robustness against outlying data values. In this discussion we will concentrate on cases that arise from the Euclidean graph, considered as a basic metric space from which we can construct a geodesic metric space usingd_α, or as the geodesic space for the metric cone extension using the d_β approach.

Note that this is a restriction of a very general approach that would start with a density f or a smooth empirical density $\hat{f}$; first, construct the geodesic metric d_α, and then, possibly apply the β-cone method to that geodesic metric space. However we shall prefer the graph approach because it leads to tractable mathematics and much faster computation; moreover, we are able to capture some of the structure of the data using the complete Euclidean graph, the Delaunay graph and the Gabriel graph.

6.1. **CAT(0), CAT(k).** As mentioned above, we restrict ourselves to the case where the base space is the Euclidean graph, i.e. $\alpha = 0$. It holds that as $\alpha \to -\infty$, all geodesics end up on the minimal spanning tree that is CAT(0). As α increases at some point, it ceases to be CAT(0) but will be CAT(k) for some k, which depends on α and indeed increases with α.

It should be stressed that the theorems on β are completely general. That is they cover arbitrary metric cones based on a geodesic metric space. If we start with the Euclidean graph as our geodesic space, it may not be CAT(0), but it can be shown that it is a CAT(k) space for some k and will eventually be CAT(0) for sufficiently small β.

6.2. **Means.** While the space is CAT(0) the intrinsic mean is unique and lies on the minimal spanning tree. For larger values of α, there may not be a unique intrinsic mean lying on the geodesic graph, but Theorem 5.2 suggests that the mean may be unique for values of the diameter that are sufficiently small relative to $k^{-\frac{1}{2}}$. For $d = 1$, at $\alpha = 0$ and $\alpha = 1$, we obtain the Euclidean intrinsic mean and the median, respectively, and these, again, may not be unique.

As noted above, the metric cone embedding the original data space becomes CAT(0) for sufficiently small β sufficiently small. The following is a possible strategy: starting with the Euclidean graph, apply the cone method decreasing β until the intrinsic mean is unique (this may happen at a larger value of β than that required to be CAT(0)). Then, project that unique mean back onto the original graph; the projection is unique by construction provided that the intrinsic mean is not at the apex of the cone. Instead of judging this method now we will wait for some empirical studies.

In the β, instead of the projection of the intrinsic mean in the cone back to the original geodesic space, we prefer the extrinsic mean in which the mean lies in the

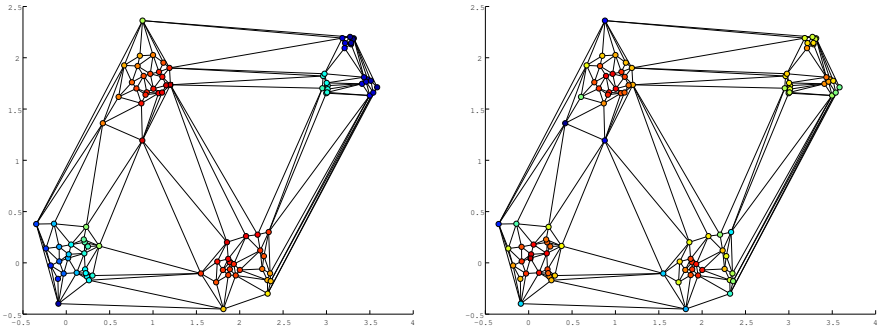

(a) $\alpha = 0, \beta = \infty$ (b) $\alpha = 0, \beta = 1$

(c) $\alpha = 1, \beta = \infty$ (d) $\alpha = 0.5, \beta = 2$

FIGURE 9. The value of $\sum_i g_\beta(d_\alpha(x_i, m))^2$ for 24 samples each from normal distributions with the s.d. $\sigma = 0.2$ and 8 samples each from normal distributions with the s.d. $\sigma = 0.05$. Initial graph: Delaunay. The value for each m is represented by the colour: red (small), blue (large), and minimum by a square.

geodesic space but, but we use the β cone metric d_β. As we decrease β we tend to get more local minima. This is the antithesis of obtaining a CAT(0) space.

Because the β metric becomes more concave as β decreases for very small β, each data point will yield a single "local mean". For reasonable values of β, each data cluster will have a local mean because the function g_β is locally linear and intra-cluster distances are small. However, when the intra-cluster distances are larger, the concavity dominates, and we obtain multiple local minima.

6.3. $\alpha > 0$ **versus** $\beta < \infty$**.** Both larger values of α and smaller values of β make the function $f(x) = \sum_i g_\beta(d_\alpha(x_i, x))^2$ less convex and can lead to multiplicity of the local means. Figures 9 shows the difference between the effects caused by positive α and finite β. Compared with the original geodesic distance (Figure 9 (a)), both (b) and (c) have multiple local minima of the function f. In (b), the three small groups of samples in the top right corner are dealt with one cluster and have a unique local minimum in their midst. In (c), each of the three small bunches has a local minimum. However, (c) tends to have local minima on a ridge of each cluster. By tuning both α and β as in (d), we can get local minima at the centre of each cluster.

6.4. **Robustness using α, β and γ.** At the beginning of Section 3.1, we showed how the univariate median can be considered as an intrinsic mean with respect to a geodesic metric, in both the population case and the empirical case. Because the univariate median is well known to be robust against outliers, we can expect means based on empirical geodesics to also be robust. In particular, the flattening of large distances when β is small gives an indication as to where to look for robust means. This flattening idea is not new. Huber [18] introduced a function similar to our g_β function, which is quadratic in a spherical neighbourhood of zero and flat outside. If the generalised mean is computed as a Fréchet mean, we should take the square root of the Huber loss function. Various generalisations, including a smoothed version of Huber loss function, have been proposed; see, for example, [16].

Figure 10 shows a geodesic graph and the value of $f(x) = \sum_i \{g_\beta(d_\alpha(x_i, x))\}^\gamma$ for 19 samples from N(0,1) and one outlier at $(0, 30)$. Figure 10 (a) uses ordinary geodesic distance and the sample closest to the outlier sample attains the minimum of f. All of $\alpha = 1$ for (b), $\beta = 1$ for (c) and $\gamma = 1$ for (d) attain the minimum of f around the mean of the normal distribution and we can see their robustness.

It is interesting that α, β and γ can control the trade-off between uniqueness of the mean and robustness of the estimation. In Euclidean space this trade-off is typically studied via the convexities of the loss function, whereas we are making the perhaps novel suggestion that the curvature is an appropriate vehicle.

7. Choosing α and β

7.1. **Choosing α.** In this section, we suggest how to select α and β empirically from the data.

First, assume that we have Euclidean data (equivalnet to $\beta = \infty$) and recall the basic effect of decreasing α from 0 to $-\infty$. At $\alpha = 0$, we make no change to the metric. As α decreases, we lose edges from the geodesic graph. That is to say from time to time, an edge that is in a particular geodesic is discarded and every geodesic that passes through that edge then has to use an alternative route.

Let us assume that at α (and under mild extra conditions), only a single edge e_0 is removed and let $d_0^{1-\alpha}$ be its length. Let $d_1^{1-\alpha}, \ldots, d_k^{1-\alpha}$ be the lengths of the edges on the new geodesic that will replace the removed edge. In addition, let there be n_0 distinct geodesics that use e_0. It is straightforward to see that *all* geodesics that use e_0 will use the new arc for an interval $[\alpha, \alpha+\epsilon)$, for sufficiently small $\epsilon > 0$. The total change in geodesic length is

$$\Delta(\alpha) = n_0(d_0^{1-\alpha} - \sum_{i=1}^{k} d_i^{1-\alpha}),$$

and it is continuous at the current α but the first derivative changes: $\Delta'(\alpha)$ is typically not zero. To see this, take the case where all the $d_i, i = 1, \ldots, k$, are equal. Then, the change in the first derivative is

$$-\frac{n_0 d_0^{1-\alpha}}{1-\alpha} \log k.$$

In graph theory, the number of geodesics using a particular edge, n_0 in our case, is sometimes called the *edge betweenness*. We might therefore refer to the term $n_0 d_0^{1-\alpha}$ as the weighted betweenness. This quantity measures changes in the

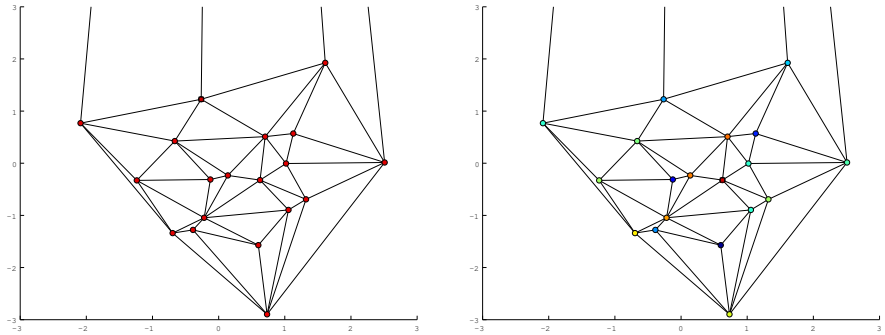

(a) $\alpha = 0, \beta = \infty, \gamma = 2$ (b) $\alpha = 1, \beta = \infty, \gamma = 2$

(c) $\alpha = 0, \beta = 5, \gamma = 2$ (d) $\alpha = 0, \beta = \infty, \gamma = 1$

FIGURE 10. Geodesic graph and value of $f(x) = \sum_i g_\beta(d_\alpha(x_i, x))^\gamma$ for 19 samples from N(0,1) and one outlier at $(0, 30)$ (far above the figure region). The value $f(x)$ for each sample point x is represented by the colours red (small) and blue (large) and the minimum is represented by a square.

configuration: if n_0 and $d_0^{1-\alpha}$ are large then a long edge with large betweenness is removed, and it is replaced by k shorter edges from the current geodesic graph.

If n_e is the betweenness of an edge e, the total betweenness of a graph $G(V, E)$ is the sum of all the individual edge betweennesses,

$$\sum_{e \in E} n_e,$$

and the weighted version is

$$\sum_{e \in E} n_e d_e^{1-\alpha},$$

which except for a scalar factor is the L_1 variance given by $\gamma = 1$, in this paper.

We shall in fact favour the use of s_1 ($\gamma = 2$), and with the above discussion in mind, we will see in Examples 1 and 2 that plots of the second derivative of $\log s_1$ do indeed have pronounced peaks and there is some matching of the α-values at the peaks with the analogous differential of the aggregate betweenness.

7.2. **Choosing β.** Section 4.1 and Figure 4 are important for understanding the β metric. We can summarise the material in a way that will indicate how to estimate

β. The first point is that β provides a metric cone. In one dimension, we wrap the real line around a circle and attach the origin. Then, the the metric cone is based on the Euclidean metric *inside the cone*. The enlarged space (referred to as the embedding space) is CAT(0) with respect to this metric.

We claim that this construction is fundamental because even in larger spaces, the geodesics are one-dimensional. Every geodesic, in some sense, has its private geodesic, but with a common vertex. Moreover, by Theorem 4.1, if the base space is CAT(0), the embedding space is CAT(0), and in both cases, we have a unique intrinsic mean and our statistics are well defined. Howeve,r if we compute the intrinsic mean restricted to the base space, e.g. Euclidean space, then the uniqueness no longer holds. As stated above, the space may not be CAT(0) for small β but may become more so for large β. We can use this to our advantage: for sufficiently large β, we expect a single maximum

$$f_\beta(m) = \sum_{i=1}^{n} g_\beta(d(x_i, m))^2,$$

but multiple maxima for smaller β, as shown in Figure 5. If we recall that the value of the function f_β for a given β is helpful in clustering, we can suggest a number of plots to show the local minima.

However, we can say more. First, note that in one dimension,

$$W_\beta(|x|) = \frac{1}{\beta}(1 - g_\beta(|x|)^2) = \frac{1}{2}\cos^2\left(\frac{\pi|x|}{2\beta}\right)$$

over $[-\beta, \beta]$ is a smooth kernel with bandwidth β. Thus, with $d(x_i, m) = |x_i - m|$, we see that

$$\frac{1}{\beta}\left(1 - \frac{f_\beta(x)}{n}\right) = \frac{1}{n}\sum_i W_\beta(|x_i - x|)$$

is a smooth density. This interpretation helps to intuitively choose β: a "typical value" of $\|x_i - x\|$, e.g. the average of $|x_i - x_j\|$, by analogy of selection of bandwidth for kernel functions.

Another option is to use cycle lengths in the geodesic graph. As can be seen from Theorem 3.4 and the proof of Theorem 4.1, if we set $\beta = |\Gamma|/2\pi$ for a cycle γ and its length $|\Gamma|$, then the metric cone generated by the cycle becomes CAT(0). The use of β is shown in Examples 1 and 3.

8. Examples

In this section, we apply the $d_{\alpha,\beta,\gamma}$ metric to real data. Because the L^γ loss function is more familiar than deformation of metrics by α and β, we will set $\gamma = 2$ and focus on α and β throughout the section.

For the d_α metric (for $\beta = \infty$, $\gamma = 2$), we briefly describe the computation. For each fixed α and each pair of points initially every $d_\alpha(x_i, x_j)$ is computed, giving a complete graph. On this graph the x_i to x_j geodesic is computed for all $i \neq j$. The geodesic graph, for this α, is then computed as the union of all such geodesics. The present version of the software computes the geodesic graph for a grid of around 100 points, depending on the range of α. As mentioned, we are interested, here, only in the range $(-\infty, 1]$ and typically consider the range $(-r, 1]$, where r is a small positive integer. For each α, we compute s_0^2 and s_1^2.

8.1. **Example 1: k-nearest neighbour classification with $d_{\alpha,\beta}$.** We apply the $d_{\alpha,\beta}$ metric to the k-nearest neighbour (k-NN) method, one of the most popular classification methods. We will see that if we can choose adequate values for α and β, the classification error can be reduced.

We use five data sets from he UCI Machine Learning Repository [4]: (i) Fisher's iris data set (number of instances $n = 150$, number of attributes $d = 4$, number of clusters $m = 3$), (ii) wine data set ($n = 178$, $d = 13$, $m = 3$), (iii) ionosphere data set ($n = 351$, $d = 32$, $m = 2$), with only real attributes, (iv) breast cancer Wisconsin (diagnostic) data set ($n = 569$, $d = 30$, $m = 2$), and (v) yeast data set ($n = 1484$, $d = 8$, $m = 10$). The average l^2 norm of each data set is normalized to be one.

The Euclidean complete graphs are used as the initial metric graphs ($\alpha = 0$), and classification is performed using the weighted k-NN method (k=10) with a common weighting $1/d^2$ where d is the distance to the neighbour data point but using $d_{\alpha,\beta}$ for various values of $\alpha \in \{-5, -4.8, \ldots, 0.8, 1\}$ and $\beta \in \{2^{-6}, 2^{-2}, \ldots, 2^5, \infty\}$. A half of the samples is selected at random as a training set and the rest half is used as a testing set to evaluate the classification result. We repeat it 1000 times and estimate the error rate.

TABLE 2. Classification by k-NN method

	k-NN with $d_{\alpha,\beta}$			with Euclidean
data set	$\hat{\alpha}$	$\hat{\beta}$	$\hat{r}$	r
(i) iris	-4.4	0.0156	**0.0334±0.0011**	0.0366±0.0011
(ii) wine	0	∞	0.2814±0.0025	0.2814±0.0025
(iii) ionosphere	-0.4	∞	0.1671±0.0018	0.1677±0.0018
(iv) cancer	0.4	2	**0.0708±0.0008**	0.0729±0.0007
(v) yeast	0.4	8	**0.4184±0.0009**	0.4227±0.0008

In Table 2, $\hat{\alpha}$ and $\hat{\beta}$ are the values attaining minimum mean classification error and $\hat{r}$ is the error rate with 95% confidence interval ($\pm 1.96(\text{std.})/\sqrt{1000}$) In addition, r is the classification error for the ordinary Euclidean k-NN. The boldfaces represent significantly smaller error rates by $d_{\hat{\alpha},\hat{\beta}}$ than Euclidean k-NN.

Figure 11 shows the geodesic graphs of the five data sets with the optimum values of α and β. To simplify the figures, 100 samples from each data set are randomly selected and the optimum values of α and β are recomputed. The shape of the sample points represents their class (we use only three types of point shapes by using the same point shape for the third and higher labeled classes for clarity of the figures). The value of $f(x) = \sum_i g_\beta(d_\alpha(x_i, x))^2$ at each sample point is represented by the different colours (red:small, blue:large). We can see that the shapes of the "optimal" geodesic graphs are variable because the optimal value of α depends on the original data spaces and the distributions.

The computation cost is linear in the number d of attributes and therefore, the number n of samples is our main concern. The heaviest part of the algorithm is to compute the shortest path length between each pair of samples. We used Floyd's algorithm [14] which requires $O(n^3)$ computations.

There is a need for a more efficient program for more than 10,000 samples. One option is to begin from the subgraph of the complete graph: for example, the union of the complete subgraph whose vertices are a subset of the samples and the

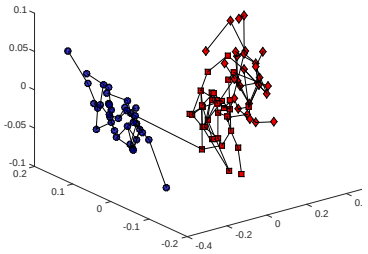

(a) iris ($\hat{\alpha} = -4.2, \hat{\beta} = 2^{-4}$)

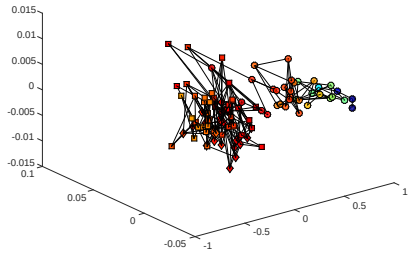

(b) wine ($\hat{\alpha} = -0.4, \hat{\beta} = \infty$)

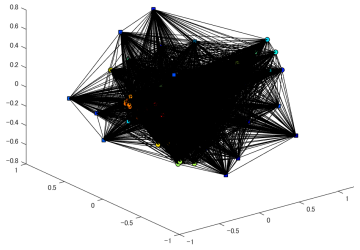

(c) ionosphere ($\hat{\alpha} = 0.8, \hat{\beta} = \infty$)

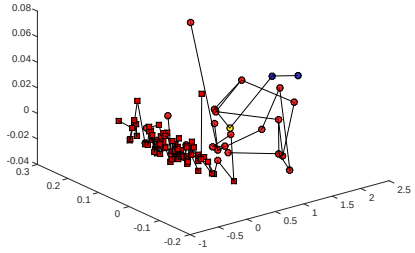

(d) cancer ($\hat{\alpha} = -3.0, \hat{\beta} = 2^{5}$)

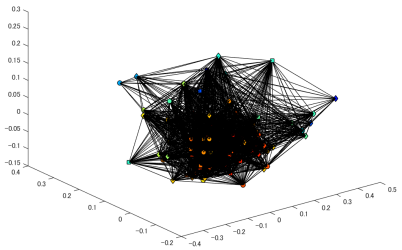

(e) yeast ($\hat{\alpha} = -0.2, \hat{\beta} = 2^{-1}$)

FIGURE 11. The geodesic graph of each data set with an optimum value of α and β.

edges connecting the remaining samples to the complete subgraph. Moreover, if we can decrease the number of edges in the geodesic graphs, Johnson's algorithm for computing the shortest path lengths can be used instead of Floyd's algorithm, because it requires only $O(|E|n + n^2 \log n)$, where $|E|$ is the number of edges.

For classification problems, we can use cross-validation to select an optimal value of α and β. However, this is not the case with clustering problems. We now extend the discussion by investigating the effect of changing α and β on the statistics of

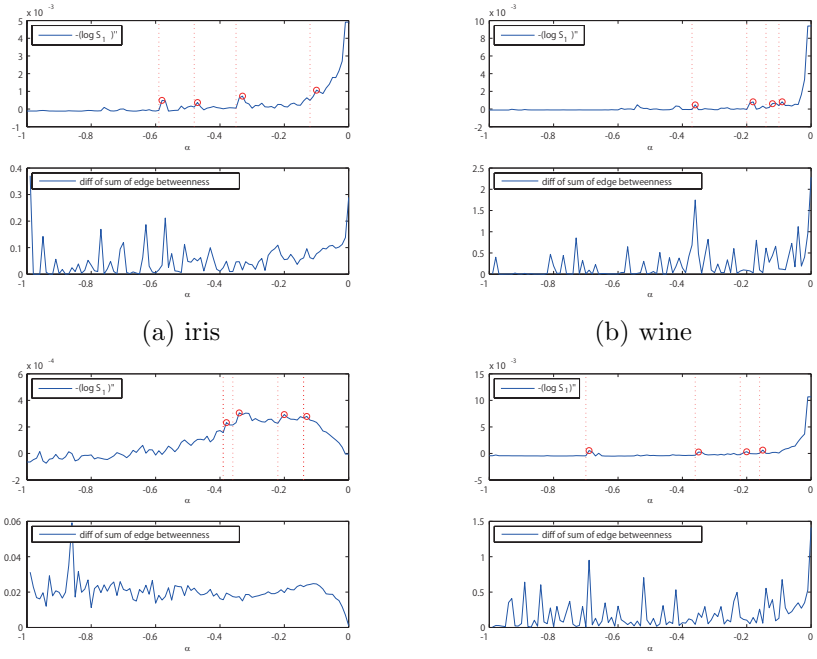

(a) iris (b) wine

(c) ionosphere (d) cancer

FIGURE 12. Upper graphs: selected peaks of $-(\log s_1)''$ (marked as circles) and values of α (dashed lines). Lower graphs: numerical differentiation of the sum of betweenness of each edge with respect to α.

s_0^2 and s_1^2 defined in the introduction. We restrict the computation to geodesics on the empirical graph discussed in Section 3.2. For the intrinsic mean, needed for s_0^2, we restrict the minimization to vertices of the graph, namely the data points. The key feature is that we first generate the graph for different values of α and then apply the β transformation to the geodesic distances on the graph: first the local metric transformation (α) and the global transformation (β).

Figures 12(a)-(d) show plots of $-(\log s_1)''$ (upper) and numerical differentiation of the total betweenness (lower) with respect to α. Note that the total betweenness is the sum of betweenness of each edge. The betweenness of each edge is defined as the number of pairs of nodes that use the edge as a part of the shortest path connecting them. Therefore, the sum of betweenness of each edge is directly affected by the topology of the geodesic graph for each value of α, and its differential can be used to detect changes in topology with respect to α. Figure 12 shows the plots to select the values of α and the summary of the selection is provided in Table 3. Figures 14 and 15 show how the values of β are selected.

8.2. **Example 2: Rainfall data.** We carry out some analysis of rainfall (precipitation) data obtained from the UK Met Office Hadley Centre[1]. Considering a single year's data we take the "dimensions" as the nine regions of the UK: South East England, South West England and Wales, Central England, North West England

TABLE 3. Selected value of α

iris	-0.59	-0.48	-0.35	-0.12
wine	-0.37	-0.20	-0.14	-0.10
ionosphere	-0.39	-0.36	-0.22	-0.14
cancer	-0.70	-0.36	-0.22	-0.16

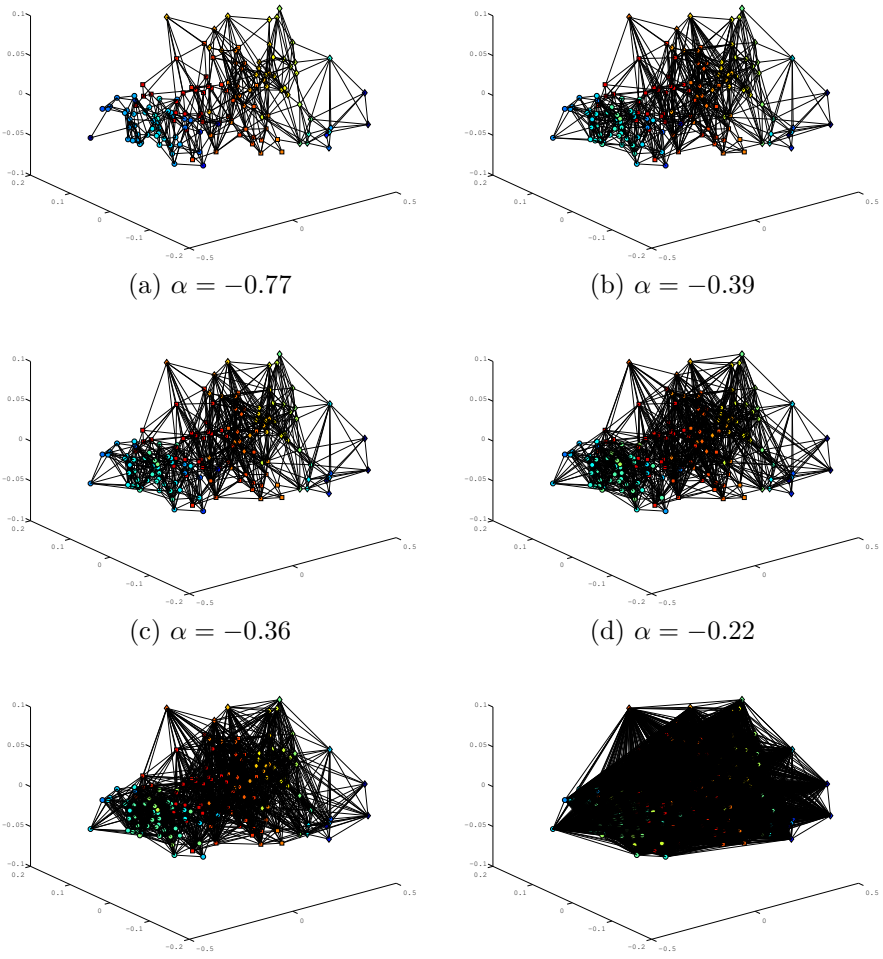

(a) $\alpha = -0.77$ (b) $\alpha = -0.39$

(c) $\alpha = -0.36$ (d) $\alpha = -0.22$

(e) $\alpha = -0.14$ (f) $\alpha = 0$

FIGURE 13. Geodesic graphs of 'iris' data for selected values of α. The data and graph are projected to a 3-dimensional subspace selected by the multidimensional scaling method.

and Wales, North East England, South Scotland, North Scotland, East Scotland and Northern Ireland, and the "points" as the 365 (or 366) days of the year. We take the years 1931 to 2014. Initially, we select values of $\alpha = 0, -0.1, -0.22, -1$ by

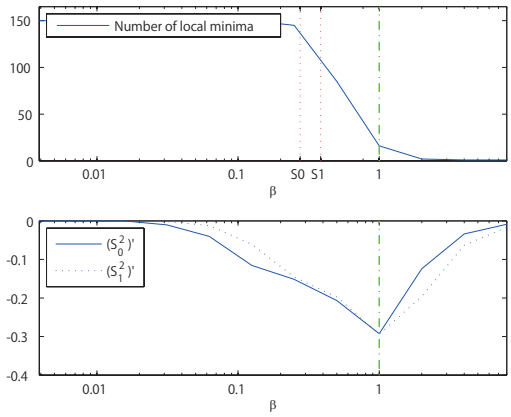

FIGURE 14. Number of local minima of function f (above) and the value of $\frac{\mathrm{d}s_0^2}{\mathrm{d}\beta}$ and $\frac{\mathrm{d}s_1^2}{\mathrm{d}\beta}$ (below) against β for iris data.

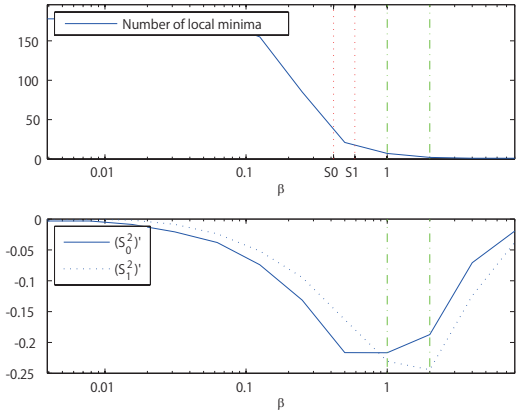

FIGURE 15. Number of local minima of function f (above) and the value of $\frac{\mathrm{d}s_0^2}{\mathrm{d}\beta}$ and $\frac{\mathrm{d}s_1^2}{\mathrm{d}\beta}$ (below) against β for wine data.

using some peaks of $-(\log s_1)''$ in Figure 16. For each year, we compute $s_0^{2/(1-\alpha)}$ for values $\alpha = 0, -0.1, -0.22, -1$. The data is presented as four time series for yearly values from 1931 to 2014 in Figure 17. The figure is consistent with an emerging consensus of increased extremes and volatility in precipitation in the UK in recent years (see Met Office, 2014, "Recent Storms Briefing").

We now discuss the choice of α. Following the discussion in Figure 16, we plot $-\frac{\partial^2}{\partial\alpha^2}\log s_1$ in the range $\alpha \in [-1, 0]$. Figure 16 shows plots for the year 2014. We select the values of α to be slightly smaller than the peaks. The local peak at approximately $\alpha = -0.22$ indicates a rapid change in the topology of the graph at

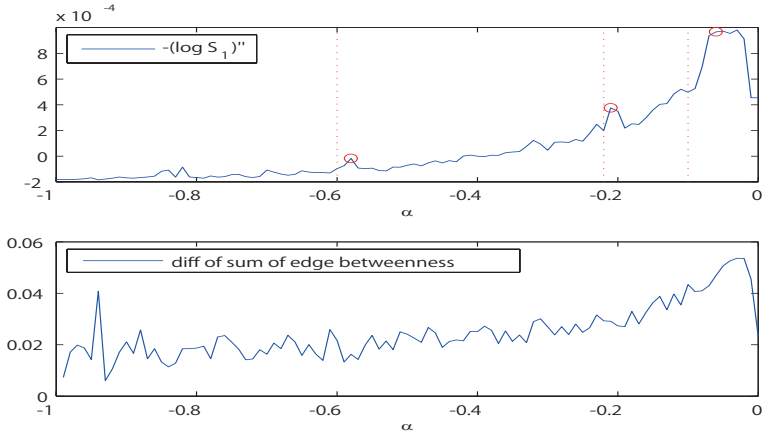

FIGURE 16. Selecting α for weather data of 2014. Upper graphs: selected peaks of $-\log s_1''$ (marked as circles) and values of α (dashed lines). Lower graphs: numerical differentiation of the sum of betweenness of each edge with respect to α.

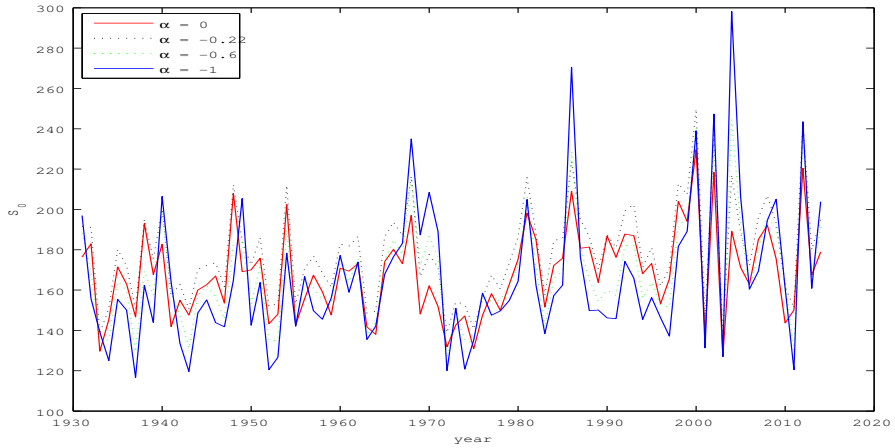

FIGURE 17. Rainfall data in UK (1931-2014): yearly time series of $s_0^{2/(\alpha-1)}$ plotted for $\alpha = 0$ (red solid line), -0.1 (green dashed line), -0.6 (black dashed line) and -1 (blue solid line).

this point: we lose a considerable number of longer edges and reveal the structure in the data as a consequence. The betweenness plot is not so revealing, except near $\alpha = -1$. The plot in Figure 18 shows on the left the distance (adjacency) matrix for the distance between points (days) on a "grey scale". Closeness can be interpreted as days that have a similar rainfall pattern over the UK; more red indicates more closeness). We see throughout the year "tartan" bands the represent periods of high and low precipitation. It is somewhat difficult to see the topology because closeness

in time does not necessarily mean closeness in weather. In the plots on the right, the days are reordered according to the standard methods employed for dendrograms, the group averaging method. See [25] and Section 8.3 for this method. We use the computed value of $d^{1-\alpha}$ in the (geodesic metric) adjacency graph, except at large distances where we take the value 10^4, for computational purposes. As α becomes more negative, we see a more tree-like structure emerging which represents the high precipitation weather at the top left and low at the other: a kind of "weather tree".

8.3. **Example 3: Sushi trees.** The work of [6] shows how to build a CAT(0) space of trees, that is where every point in the space represents a tree: see also [24]. In this example, we start with such a space and apply the clustering method based on the β-metric, d_β. In statistics, tree data is often generated from data via a different structure with post-processing. This is true for examples in phylogenetics, where trees are constructed from genome data. In our example, we start with a data set of preferences for sushi, based on individual ranking: see [19]. We used the data set in which individuals rank 10 sushi dishes: shrimp, sea eel, tuna, squid, sea urchin, salmon roe, egg, fatty tuna, tuna roll and cucumber roll.

First from each individual ranking we make a preference distance matrix for the sushi dishes by the difference in their ranks. To obtain a single tree, we sample 20 individuals without replacement, and from their average distance matrix, we construct a tree using the group averaging method for dendograms see [25]. The tree is a metric tree in that the method gives weighted edges. We select 100 samples of size 20 for women and the same for men. This gives, therefore, 100 trees for men and 100 for women. Each tree is a point in CAT(0) space and we can compute the geodesic distance $d(x_i, x_j)$ between each pair of trees x_i, x_j. We employ a version of the algorithm described in [29].

We could compute the theoretical intrinsic mean, but for computational purposes we take the minimum of

$$f_{\alpha,\beta}(m) = \sum_i g_\beta(d_\alpha(x_i, m))^2,$$

where m ranges over the trees in the sample. When $\beta = \infty$, the function g_β is the identity and we obtain a unique mean by the CAT(0) property, the uniqueness typically extends to the restriction of m to sample trees. However, for finite β, this uniqueness can break down, but as explained, we can use this property for clustering.

We find it interesting that the the use of $d_{\alpha,\beta}$ with $\alpha = 0$ is not effective for this tree data. However, by applying negative values of α after inspection of the α, similar to Example 2, we selected the value of $\alpha = -4$ and it was effective in separating the samples of men and women.

Applying the method to all 200 trees, we select a value of $\beta = 1.1$ at which there are two distinct minima of $f_{\alpha\beta}(m)$, that is two trees that may be representative of two clusters. Note that we define a local minima, as the minimum over a moving window of size 1 at each data point (tree). Table 4 gives the geodesic distances between the cluster centres and the intrinsic means for the dendrograms of men and women.

Figure 19 shows the trees (dendrograms) for men and women: the first two are the intrinsic means and the next two are centres selected by the minima of $f_{\alpha,\beta}(m)$, for the selected values of α and β. Figure 20 was constructed by first finding the two

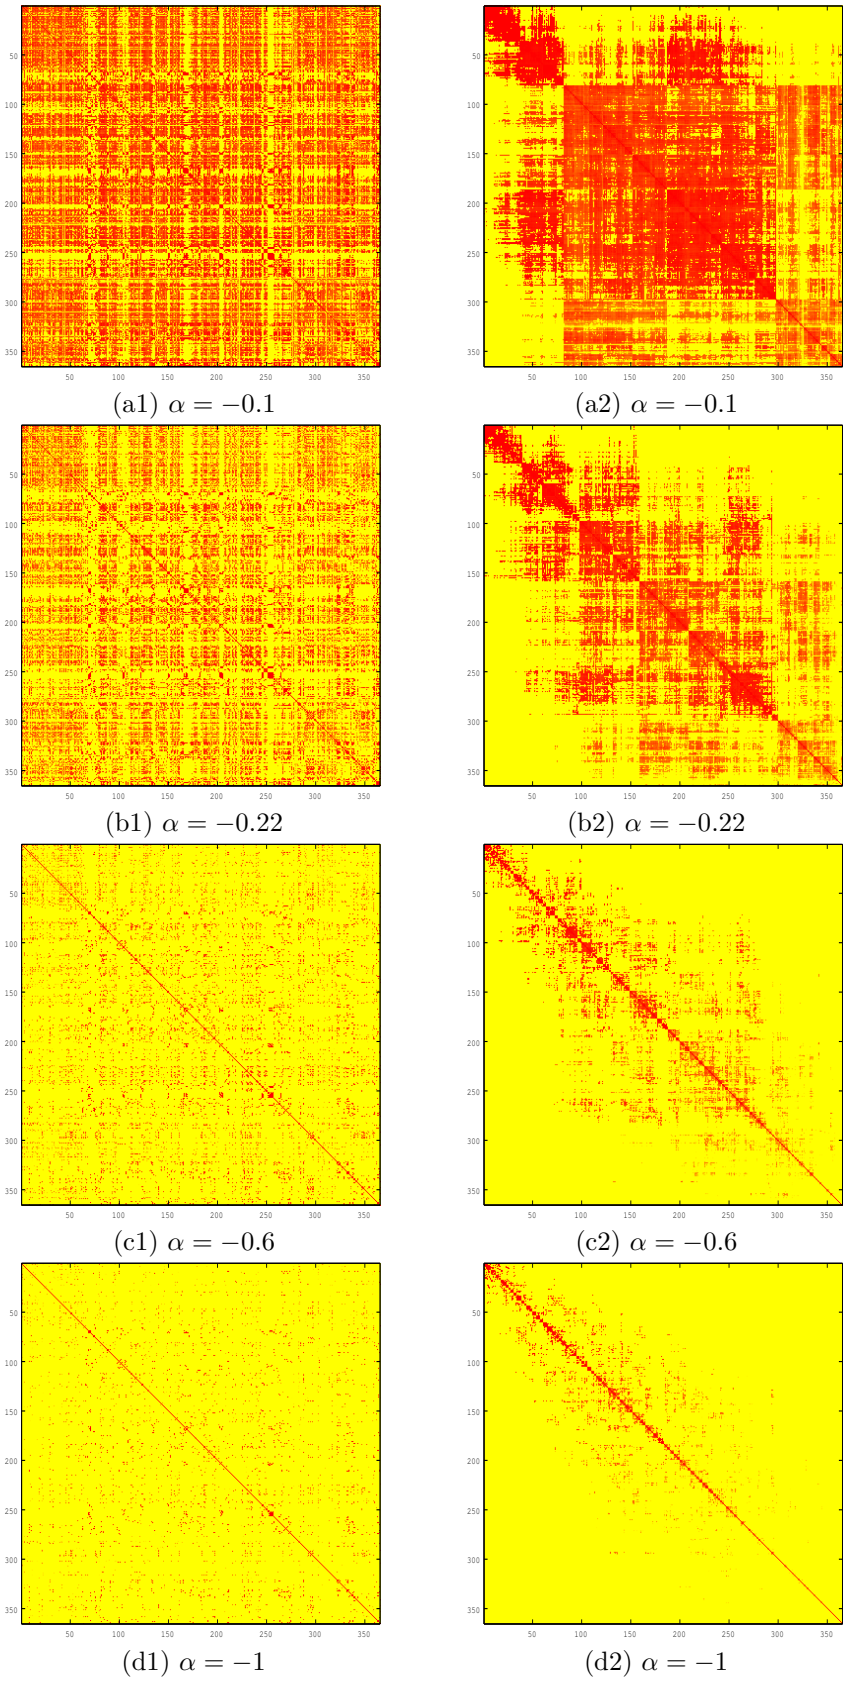

(a1) $\alpha = -0.1$ (a2) $\alpha = -0.1$

(b1) $\alpha = -0.22$ (b2) $\alpha = -0.22$

(c1) $\alpha = -0.6$ (c2) $\alpha = -0.6$

(d1) $\alpha = -1$ (d2) $\alpha = -1$

FIGURE 18. Raw and reordered adjacency matrices for $\alpha = 0, -0.1, -0.6$ and -1.

TABLE 4. Geodesic distances between the cluster centres and the intrinsic means for the dendrograms of men and women

	Men (a)	Women (b)
Cluster 0 (c)	0	1.13
Cluster 1 (d)	1.09	0.79

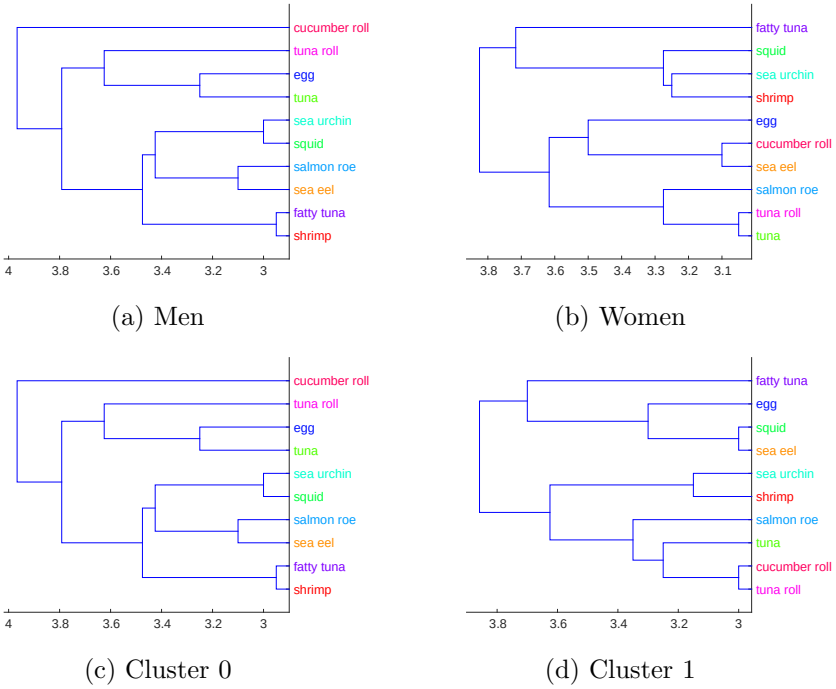

(a) Men (b) Women

(c) Cluster 0 (d) Cluster 1

FIGURE 19. Dendrograms of intrinsic means of (a) men and (b) women dendrograms. (c) and (d) are two cluster centre dendrograms selected by the local minima of $f_{\alpha,\beta,\gamma}$.

local minima of $f_{\alpha,\beta}(m)$ and then plotting the values of $f_{\alpha,\beta}(m)$ along a geodesic of maximal length passing through the two minima. It transpires that there are only six trees on this geodesic (including the two special points) but we are able to interpolate between the points by extending the range of the argument m to include points in the edge, considered as a subset of the metric graph.

9. CONCLUSION

The α metric is a deformation of the starting geodesic, and as α becomes more negative, the geodesic graph, namely the union of all the geodesics, becomes sparser, and in our graph representation, it becomes a tree. The space is CAT(k) with smaller k and finally becomes a tree, at which point the space becomes CAT(0). It is quite difficult to see the tree computation because of the numerous short edges, but for moderate values of α such as -1, the structure is tree-like. Abrupt changes

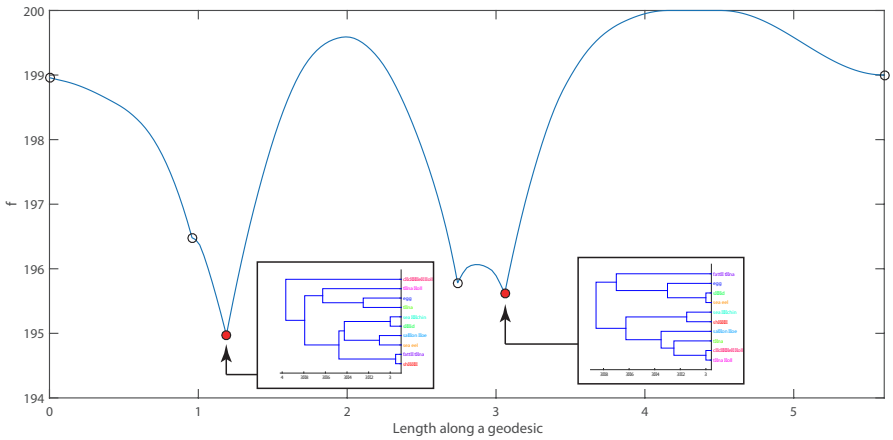

FIGURE 20. Graph of $f_{\alpha,\beta,\gamma}$ along a geodesic passing its two local minima ($\alpha = -4, \beta = 1.1, \gamma = 2$). The value at each sample point on the geodesic is marked as a circle. The two red circles correspond to the local minima.

in various statistics as α changes reveal topological changes in the structure of the geodesic graph, a fact that can be used to tune α.

The β metric is "non-geodesic" because although the function g_β operates on a geodesic, that does not mean that the space is a geodesic space in the formal sense. However the cone construction yields a geodesic metric space, which is CAT(k) with a lower value of k than the original space, and indeed may be CAT(0). If the β metric is projected back to the original space, that space can have a non-convex f-function with larger k. This is useful for finding clusters because of multiple minima of the f-function, which is itself similar to a kernel. The means obtained by the β-metric may represent the first study of an extrinsic mean via embedding in non-Euclidean spaces and the first application of metric cones to statistics and data analysis.

We believe that the curvature of the data space underlying this work demands further investigation whereby connections should be established with recent developments related to empirical geodesic graphs, for example in manifold learning. One important direction should be the effect of the curvature of the space on the trade-off between the uniqueness of the Frechét means and the robustness of estimation. To this end, α, β and γ can be seen as parameters that can be tuned to change the curvature and hence study the trade-off.

APPENDIX A.

A.1. **Delaunay and Gabriel graphs.** The Delaunay complex is usually defined via the Voronoi diagram. Let $\{x_i, i = 1, \ldots, n\}$ be distinct points in $\mathbb{R}^d$. Define for $i \neq j,\ i, j = 1, \ldots, n$ the $d-1$ hyperplane A_{ij} perpendicular to and bisecting the line from x_i to x_j. Again for $i \neq j$ let H_{ij} be the half space bounded by A_{ij} and

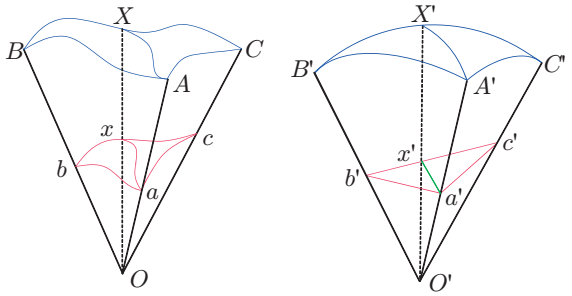

FIGURE 21. The cone spanned by geodesics $\overline{ab}$, $\overline{ac}$, $\overline{bc}$ and $\overline{ax}$ (left) and the cone spanned by a comparison triangle $\Delta a'b'c'$ (right).

containing x_i. Then, the Voronoi cell of x_i is

$$V_i = \bigcap_{j \neq i} H_{ij}.$$

The elementary definition is that V_i is all those no further from x_i than any other x_j. The Delaunay complex is a simplicial complex with vertices at x_i and one-dimensional edges between x_i and x_j if and only if H_{ij} is a face of V_i (and V_j). The vertices x_i and the edge set then define the Delaunay graph: the graph of the complex. Note that in the text we shall assume a general position when no more than $d+1$ Voronoi cells can have a common point, and this guarantees that the Delaunay complex has maximum simplexes of dimension d and $d+1$ vertices

A well-known property of the Voronoi construction, in the general position case, is that the circumscribed sphere of a Delaunay simplex contains only the $d+1$ vertices of the simplex, and in fact this is a defining property of the maximal simplexes. The Gabriel graph is defined in an intuitively similar way: the edge from x_i to x_j is included if and only if the unique sphere whose diameter is that edge contains no third point x_k. It is straightforward to show that the Gabriel graph is a subgraph of the Delaunay graph. See [28] [26] [2] for theory, algorithms and applications of Delaunay and Gabriel graphs. The Delaunay complex has come to prominence in the areas of discrete tube theory and persistent homology [12].

A.2. **Proof of Theorem 4.1.** (1) Denote the mapped points of a, b, c and x by the projection $\tilde{\mathcal{X}}_\beta \to \mathcal{X}_\beta$ as A, B, C and X, respectively, as shown in Figure 21 (left). Denote the origin of the metric cone as O. If the sum of the lengths of the geodesics $\widetilde{AB}$, $\widetilde{AC}$ and $\widetilde{BC}$ in $\mathcal{X}_\beta$ exceeds $2\pi\beta$, it is easy to see that the cone spanned by $\overline{ab} \cup \overline{ac} \cup \overline{bc}$ becomes CAT(0) and Δabc satisfies the CAT(0) property. Therefore, assume that $|\widetilde{AB}| + |\widetilde{AC}| + |\widetilde{BC}| \leq 2\pi\beta$.

Next, let $\Delta a'b'c'$ be a comparison triangle of Δabc and let x' be a point on a geodesic $\overline{b'c'}$ such that $|\overline{bx}| = |\overline{b'x'}|$. Thus, $|\overline{a'x'}| < |\overline{ax}|$. Arrange the points a', b' and c' in a three-dimensional Euclidean space with origin O' such that the lengths of $\overline{O'a'}$, $\overline{O'b'}$ and $\overline{O'c'}$ are equal to the lengths of $\overline{Oa}$, $\overline{Ob}$ and $\overline{Oc}$, respectively. Denote the radial projection of a', b', c' and x' to a unit sphere as A', B', C' and X', respectively, as shown in Figure 21 (right). By the definition of a metric cone,

$|\overline{Ox}| = |\overline{O'x'}|$ and the geodesics $\widetilde{A'B'}$, $\widetilde{A'C'}$, $\widetilde{B'C'}$ and $\widetilde{A'X'}$ in the unit sphere are arcs satisfying $|\widetilde{A'B'}| = |\widetilde{AB}|$, $|\widetilde{A'C'}| = |\widetilde{AC}|$, $|\widetilde{B'C'}| = |\widetilde{BC}|$ and $|\widetilde{B'X'}| = |\widetilde{BX}|$.

From the argument above, $|\widetilde{A'B'}| + |\widetilde{A'C'}| + |\widetilde{B'C'}| = |\widetilde{AB}| + |\widetilde{AC}| + |\widetilde{BC}| \leq 2\pi$. Since the unit sphere has a positive constant curvature and $\mathcal{X}_\beta$ is CAT(0), $|\widetilde{A'X'}| > |\widetilde{AX}|$. However, since $|\overline{Oa}| = |\overline{O'a'}|$ and $|\overline{Ox}| = |\overline{O'x'}|$, $|\widetilde{A'X'}| > |\widetilde{AX}|$ implies that $|\widetilde{a'x'}| > |\widetilde{ax}|$ by the property of a metric cone. Thus, Δabc has CAT(0) property and (1) of the theorem is proved.

(2) Assume that $0 < \beta_1 < \beta_2 < \infty$ and a metric cone $\tilde{\mathcal{X}}_{\beta_1}$ is not CAT(0) for proving the latter half of the theorem by contradiction. Then, there is a geodesic triangle $\Delta a_1 b_1 c_1$ in $\tilde{\mathcal{X}}_{\beta_1}$ and a point x_1 on the geodesic $\overline{b_1 c_1}$ such that the geodesic $\overline{a_1 x_1}$ is longer than the corresponding geodesic of a comparison triangle. By defining $A_1, B_1, C_1, X_1, a_1', b_1', c_1', x_1', A_1', B_1', C_1'$ and X_1' as above, we can say that $|\widetilde{A_1X_1}| > |\widetilde{A_1'X_1'}|$.

Next, each of A_1, B_1, C_1 and X_1 corresponds to a point in $\mathcal{X}_{\beta_1}$ and we can consider the corresponding points A_2, B_2, C_2 and X_2 in the other metric cone $\tilde{\mathcal{X}}_{\beta_2}$. When restricted to $\mathcal{X}_{\beta_1}$, a geodesic $\widetilde{A_1X_1}$ is just a rescaling of $\widetilde{A_2X_2}$ and $|\widetilde{A_1X_1}| = \frac{\beta_1}{\beta_2}|\widetilde{A_2X_2}|$.

Now, $\Delta A_2' B_2' C_2'$ is a geodesic triangle on the unit sphere, but after rescaling by $\frac{\beta_2}{\beta_1}$, we can get a geodesic triangle $\Delta A_2'' B_2'' C_2''$ on a sphere of radius $\frac{\beta_2}{\beta_1}$ whose edges have the same length as $\Delta A_1' B_1' C_1'$. By a known result on spherical triangles with the same edge lengths on different spheres, a larger radius implies a "thinner" triangle and $|\widetilde{A_1X_1}| < |\widetilde{A_2''X_2''}|$ where X_2'' is a point on the geodesic $\widetilde{B_2''C_2''}$ such that $|\widetilde{B_1X_1}| = |\widetilde{B_2''X_2''}|$.

Combining all the arguments gives

$$|\widetilde{A_2X_2}| = \tfrac{\beta_2}{\beta_1}|\widetilde{A_1X_1}| > \tfrac{\beta_2}{\beta_1}|\widetilde{A_1'X_1'}| > \tfrac{\beta_2}{\beta_1}|\widetilde{A_2''X_2''}| = |\widetilde{A_2'X_2'}|.$$

Select a non-degenerate geodesic triangle in $\tilde{\mathcal{X}}_{\beta_2}$ by selecting arbitrary points a_2, b_2 and c_2 on the geodesics $\overline{OA_2}$, $\overline{OB_2}$ and $\overline{OC_2}$ in $\tilde{\mathcal{X}}_{\beta_2}$, respectively, and let x_2 be the intersection point of $\overline{OX_2}$ and $\overline{b_2c_2}$. Then, by $|\widetilde{A_2X_2}| > |\widetilde{A_2'X_2'}|$, we can say that $|\overline{a_2x_2}| > |\overline{a_2'x_2'}|$. This implies that $\tilde{\mathcal{X}}_{\beta_2}$ is not CAT(0) and (2) of the theorem is proved.

(3) For $k = 0$, the statement holds by (1). For $k > 0$ and $\beta \leq \pi$, it is sufficient to prove for $\beta = \pi/\sqrt{k}$ by (2). Let Δabc be a geodesic triangle in $\tilde{\mathcal{X}}_\beta$ and let ΔABC be a geodesic triangle in $\mathcal{X}_\beta$. Let A, B, C be the projection of a, b, c, respectively. If the perimeter of ΔABC is longer than or equal to 2π, the cone spanned by the perimeter becomes CAT(0) by the same argument as that for (1). Therefore, Δabc is CAT(0) and satisfies the CAT(0) property.

If the perimeter of ΔABC is smaller than 2π, since $\mathcal{X}$ is CAT(k) and $\mathcal{X}_\beta$ is CAT(1), for any $X \in \widetilde{BC}$, $\widetilde{BX}$ is shorter than the corresponding great arc $\widetilde{B'X'}$ of a comparison triangle $\Delta A'B'C'$, which is a spherical triangle on the unit sphere. Since a comparison triangle $\Delta a'b'c'$ of Δabc can be embedded on the cone spanned by $\Delta A'B'C'$, $\widetilde{bx}$ is shorter than the corresponding line segment $\widetilde{b'x'}$. This means the Δabc satisfies the CAT(0) property. □

A.3. **Proof of Theorem 3.5.** The proof is by contradiction. We need to show that if we take an edge of the complete graph that is *not* a edge of the Delaunay graph then it can be replaced by a shorter edge that is in the Delaunay graph. Let e_{ij} be our test edge from vertex i to vertex j. The edge cannot be a boundary edge of the convex hull of the vertices, because it is well known that every such edge is in the Delaunay graph. Thus, one of vertices i or j must be an interior point, say i. Consider the edge e_{ij} as it leaves i. Since i is completely surrounded by Delaunay cells that have i as a vertex it must pass through one such cell that we call D.

Replace one vertex other than i, say i' of D, by the vertex j. Now, we still have $n+1$ vertices for $D^* = \{D \backslash i'\} \cup j$ and we can construct their (unique) circumscribing sphere S. Now, Delaunay cells are uniquely determined as those sets of $n+1$ vertices whose circumscribing sphere does not contain any other points. From this we see that i' cannot be out side S, because otherwise its circumscribing sphere would contain j. Our claim is that any $n-1$-dimensional hyperplane H that contains the edge e_{ij} has at least one edge of $\{D \setminus i\}$, say e_{ir}, that must be on the "small side" of H in relation to S; that is the side not containing the centre of S. Any such edge, and therefore, e_{ir} must be shorter than e_{ij}. Then, if e_{ij} were in the minimal spanning tree, we could replace it by e_{ir}, giving a contradiction.

The proof for the Gabriel graph is more straightforward. We need to show that the minimal spanning tree has the Gabriel property. Take any edge e_{ij} of the tree if the diametric sphere contains the third vertex k then we can replace e_{ij} by e_{ik} or e_{jk}, both of which have a shorter "length" than e_{ij}. This gives a contradiction.

A.4. **Proof of Theorem 5.2.** (1) Although this is a known result, for example [23] [13], we show a short proof. Since a comparison triangle for the CAT(k) property is on a sphere of radius $1/\sqrt{k}$, first consider the unit sphere S^2 and the geodesic distance d on it. Take three points $a, b, c \in S^2$ and think of the convexity of $d(x, a)$ for $x \in \widetilde{bc}$. Without losing generality, assume that a is on the plane $y = 0$ and $\widetilde{bc}$ is on the plane $z = 0$ and let $b = (\cos\theta_0, \sin\theta_0, 0)$, $x = (\cos\theta, \sin\theta, 0)$ and $a = (\cos\psi, 0, \sin\psi)$ for $\theta_0, \theta \in (-\pi, \pi]$,$\psi \in [-\pi/2, \pi/2]$.

Thus, $d(a, x) = \arccos(a^\top x) = \arccos(\cos\theta\cos\psi)$ and $d(b, x) = |\theta - \theta_0|$. Note that for $\psi = 0$, $d(a, x) = |\theta|$ for $\psi = 0$ is a convex of $d(b, x)$. For a $\psi \neq 0$, $d(a, x) \leq \pi/2$ for $x \in \widetilde{bc}$, $d(a, x)$ is a convex of $d(b, x)$ iff $a^\top x = \cos\theta\cos\psi \geq 0$ since $\frac{\partial^2}{\partial\theta^2} d(a, x) = \cos(\theta)\cos(\psi)\sin^2(\psi)(1 - \cos^2\theta\cos^2\psi)^{-2/3}$. This means that if $d(a, x) \leq \pi/2$ for $x \in \widetilde{bc}$, $d(a, x)$ is a convex of $d(b, x)$.

If $\mathcal{X}$ is CAT(k) and has a diameter of at most $\pi/(2\sqrt{k})$, there is a comparison triangle $\Delta a'b'c'$ on a sphere of radius $1/\sqrt{k}$ such that its perimeter is at most $3\pi/(2\sqrt{k})$ and $d(a', x')$ is a convex of $d(b', x')$ for each $x' \in \widetilde{b'c'}$ because of the argument above after scaling by $1/\sqrt{k}$.

(2) is well known. See [13].

(3) We show an example of the probability measure with a three-point support on S^2 such that the diameter is larger than $\pi/2$ but can be arbitrarily close to $\pi/2$ and the uniqueness of the intrinsic L_1-mean fails.

Take $Q = (1, 0, 0)$, $P_1 = (-\sin\theta, \cos\theta\sin\psi, \cos\theta\cos\psi)$, $P_2 = (-\sin\theta, \cos\theta\sin\psi, \cos\theta\cos\psi)$ and $X = (-\sin\theta, 0, \cos\theta)$ with $\theta, \psi \in (0, \pi/2)$, as in Figure 22. Let $Y = (-\sin\theta', 0, \cos\theta')$ be the mid point of $\widetilde{P_1P_2}$ for $\theta' > \theta$. Put the point masses m_1 at Q and M at P_1 and P_2, and assume that there is a unique intrinsic median μ.

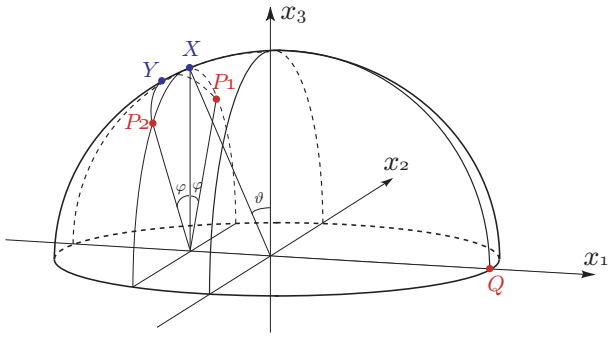

FIGURE 22. P_1,P_2,Q, X and Y on the unit hemisphere.

By the symmetry, μ must be on the arc $\widetilde{QY}$, and if we change the ratio M/m, μ moves continuously on $\widetilde{QY}$. Thus, we can set $\mu = X$ by tuning M/m adequately. However, L_1-dispersion from X becomes $S_X = md(\theta, X) + Md(P_1, X) + Md(P_2, X) = m(\pi/2 + \theta) + 2M\psi$, and L_1-dispersion from P_1 becomes $S_{P_1} = md(\theta, P_1) + Md(P_2, P_1) = m(\pi/2 + \theta) + 2M\psi$. This contradicts the assumption of X being the unique L_1-intrinsic mean. Since we can set θ and ψ as arbitrarily small positive numbers, $D_{L_1} \le \pi/(2\sqrt{k})$. However, by (1), $D_{L_1} \ge D_{\text{convex}} \ge \pi/(2\sqrt{k})$; thus, $D_{L_1} = \pi/(2\sqrt{k})$. □

A.5. **An upper bound of D_{L_γ}.**

Theorem A.1. *If $\mathcal{X}$ is a surface with a constant curvature $k > 0$,*

$$D_{L_\gamma} \le \frac{1}{\sqrt{k}} \left(\theta_0(\gamma) + \frac{\pi}{2} \right)$$

where $\theta_0(\gamma)$ is the inverse function of

$$\gamma_0(\theta) = \begin{cases} \left[\log_2 \frac{(\pi + 2\theta)}{2 \arccos\left\{ \sin^2\theta + \cos^2\theta (1 - 1/2(1 - \sin\theta))^{1/2} \right\}} \right]^{-1} & \text{for } 0 \le \theta \le \pi/6, \\ \left[\log_2 \frac{\pi - 2\theta}{\arccos(\sin^2\theta)} \right]^{-1} & \text{for } \pi/6 < \theta \le \pi/2 \end{cases}$$

for $1 \le \gamma < 2$ and $\theta_0(\gamma) = \pi/2$ for $\gamma \ge 2$.

The graph of $\theta_0(\gamma)$ is shown in Figure 8.

Proof. The proof is similar to that of Theorem 5.2(3). We consider two cases of arrangement of three points Q, P_1 and P_2.

C1: $Q = (1, 0, 0)$, $P_1 = (-\sin\theta, \cos\theta \sin\psi, \cos\theta \cos\psi)$, $P_2 = (-\sin\theta, -\cos\theta \sin\psi, \cos\theta \cos\psi)$ where $\psi = \arccos\left\{ \left(1 - \frac{1}{2(1 - \sin\theta)} \right)^{1/2} \right\}$, as shown in Figure 22. This satisfies $|\widetilde{P_1Q}| = |\widetilde{P_1P_2}|$.

C2: $Q = (1, 0, 0)$, $P_1 = (-\sin\theta, \cos\theta, 0)$, and $P_2 = (-\sin\theta, -\cos\theta, 0)$, as shown in Figure 23.

We put point masses m at Q and M at P_1 and P_2.

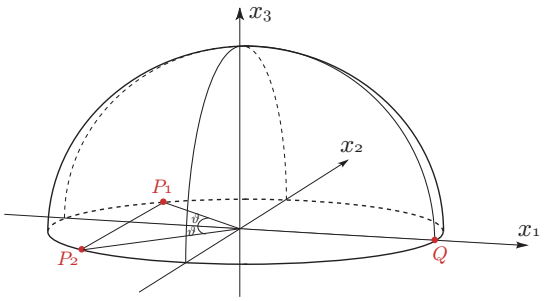

FIGURE 23. P_1,P_2 and Q on the unit hemisphere.

For $1 \leq \gamma \leq \left\{\log_2\left(\frac{2\pi}{3\arccos(1/4)}\right)\right\}^{-1}$, we consider C1. As in the proof of Theorem 5.2(3), we can set $\mu = X = (-\sin\theta, 0, \cos\theta)$. Let S_X and S_{P_1} denote L_γ-dispersion from X and P_1, respectively. Then,

$$S_X = m(\pi/2+\theta)^\gamma + 2M\{\arccos(1-\cos^2\theta(1-\cos\psi))\},$$
$$S_{P_1} = m(\pi/2+\theta)^\gamma + M\{\arccos(1-2\cos^2\theta\sin^2\psi)\}.$$

Therefore, $S_X < S_{P_1}$ is equivalent to

$$\arccos(1-2\cos^2\theta\sin^2\psi) < 2^{1/\gamma}\arccos\{1-\cos^2\theta(1-\cos\psi)\}.$$

By setting $\psi = \arccos\{(1-1/2(1-\sin\theta))^{1/2}\}$, this is equivalent to $\gamma < \gamma_0(\theta)$ and also $\theta < \theta_0(\gamma)$. Thus, if we set $\theta \geq \theta_0(\gamma)$, C1 becomes an example of a non-unique intrinsic L_1-mean of diameter $\theta + \pi/2$.

For C2, $S_X < S_{P_1}$ is equivalent to $\pi - 2\theta < 2^{1/\gamma}\arccos(\sin^2\theta)$, and we can prove that it becomes a similar example. After scaling by $1/\sqrt{k}$, these examples give the upper bound on D_{L_γ}. □

REFERENCES

1. Alexander, L.V. and Jones, P.D., Updated precipitation series for the U.K. and discussion of recent extremes, Atmospheric Science Letters, vol.1, 2001, doi:10.1006/asle.2001.0025.
2. Aurenhammer, F., Voronoi diagrams–a survey of a fundamental geometric data structure, *ACM Computing Surveys*, vol.23, no.3, pp. 345-405, 1991.
3. Baccelli, F., Tchoumatchenko, K. and Zuyev, S., Markov paths on the Poisson-Delaunay graph with applications to routeing in mobile networks, *Advances in Applied Probability*, pp. 1–18, 2000.
4. Bache, K. and Lichman, M., UCI Machine Learning Repository [http://archive.ics.uci.edu/ml]. Irvine, CA: University of California, School of Information and Computer Science, 2013.
5. Bhattacharya, A., and Bhattacharya, R., *Nonparametric inference on manifolds: with applications to shape spaces*, vol.2. Cambridge University Press, 2012.
6. Billera, L. J., Holmes, S. P. and Vogtmann, K., Geometry of the space of phylogenetic trees, *Advances in Applied Mathematics*, vol.27, no.4 pp. 733–767, 2001
7. Bose, P. Devroye, L., Löffler, M., Snoeyink, J. and Verma, V., Almost all Delaunay triangulations have stretch factor greater than $\pi/2$, *Computational Geometry*, vol.44, no.2, pp. 121–127, 2011.
8. Cignoni, P., Montani, C. and Scopigno, R. DeWall: A fast divide and conquer Delaunay triangulation algorithm in Ed, *Computer-Aided Design*, vol.30, no.5, pp. 333–341, 1998.

9. Das, G., Heffernan, P. and Narasimhan, G., Optimally sparse spanners in 3-dimensional Euclidean space, *Proceedings of the ninth annual symposium on Computational geometry*, pp. 53–62, 1993.
10. de Berg, M., Cheong, O., van Kreveld, M. and Overmars M., *Computational Geometry: Algorithms and Applications*, Springer-Verlag, 2008.
11. Deza, M. M. and Deza, E., *Encyclopedia of distances*, Springer Berlin Heidelberg, 2009.
12. Edelsbrunner, H. and Harer J. L., *Computational Topology*, AMS, 2009
13. Espínola, R. and Fernández-León, A., CAT (k)-spaces, weak convergence and fixed points, *Journal of Mathematical Analysis and Applications*, vol.353, no.1, pp. 410–427, 2009.
14. Floyd, R., Algorithm 97: Shortest Path, *Communications of the ACM*, vol.5, no.6, p. 345, 1962.
15. Gromov, M., *Hyperbolic groups*, Springer New York, 1987.
16. Hampel, F. R., Ronchetti, E. M., Rousseeuw, P. J. and Stahel, W. A., Robust statistics: the approach based on influence functions, vol.114, Wiley, 2011.
17. Hotz, T., Extrinsic vs Intrinsic Means on the Circle *Geometric Science of Information Lecture Notes in Computer Science*, vol.8085, pp. 433–440, 2013,
18. Huber, P. J., Robust Estimation of a Location Parameter, *Annals of Statistics*, vol.53, pp. 73–101, 1964.
19. Kamishima, T Nantonac Collaborative Filtering: Recommendation Based on Order Responses", Knowledge Discovery and Data mining: KDD2003. 583–588 (2003)
20. Keil, J. M. and Gutwin, C. A., Classes of graphs which approximate the complete Euclidean graph, *Discrete & Computational Geometry*, vol.7, no.1, pp. 13–28, 1992.
21. Kendall, D. G., Barden, D., Carne, T. K. and Le, H., Shape and shape theory, vol.500, John Wiley & Sons, 2009.
22. Kendall, M. G. and Moran, P. A. P. *Geometrical Probability*. Griffin, London, 1963.
23. Kendall, W. S., Probability, convexity, and harmonic maps with small image I: uniqueness and fine existence, *Proceedings of the London Mathematical Society*, vol.3, no.2, pp. 371-406, 1990.
24. Kobayashi, K. and Orita, M. Permutation test for dendrograms and its application to the analysis of mental lexicons. http://arxiv.org/abs/1403.2845, 2014.
25. Lance, G. N. and Williams, W. T. A general theory of classificatory sorting strategies. I. Hierarchical systems. Computer Journal, Vol.9, pp.373–80, 1967.
26. Matula, D. W. and Robert R. S., Properties of Gabriel graphs relevant to geographic variation research and the clustering of points in the plane, *Geographical analysis*, vol.12, no.3 pp. 205–222, 1980.
27. Miller, E., Owen, M. and Provan, J. S., Averaging metric phylogenetic trees, *arXiv preprint*, arXiv:1211.7046, 2012.
28. Okabe, A., Boots, B., Sugihara, K. and Chiu, S. N., *Spatial tessellations: concepts and applications of Voronoi diagrams*, 2nd. ed., Wiley, 2000.
29. Owen, M. and Provan J. S., A fast algorithm for computing geodesic distances in tree space, *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, vol.8, no.1, pp. 2–13, 2011.
30. Xia, G., Improved upper bound on the stretch factor of delaunay triangulations, *Proceedings of the 27th annual ACM symposium on Computational geometry*, pp. 264–273, 2011.

Institute of Statistical Mathematics, Tokyo
E-mail address: kei@ism.ac.jp

London School of Economics
E-mail address: h.wynn@lse.ac.uk